

Detecting Superfluids, Exciting the Higgs Mode and Enhanced Cooling of Dimers in the BEC-BCS Crossover

Dissertation

zur

Erlangung des Doktorgrades (Dr. rer. nat.)

der

Mathematisch-Naturwissenschaftlichen Fakultät

der

Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Daniel Eberz

aus

Troisdorf

Bonn 2024

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

Gutachter / Betreuer: Prof. Dr. Michael Köhl
Gutachter: Prof. Dr. Sebastian Hofferberth

Tag der Promotion: 10.04.2025
Erscheinungsjahr: 2025

Abstract

In this thesis, the BEC-BCS crossover is experimentally investigated using a quantum simulator apparatus. We prepare a degenerate, interacting fermionic sample by cooling atoms in two of the lowest hyperfine states of ^6Li in a crossed optical dipole trap. Interactions between the two states are controlled by means of a broad magnetic Feshbach resonance, and we adjust the samples' temperature and density by precisely tuning the trapping potential. This setup allows us to access and probe the entire BEC-BCS crossover.

A key property of the BEC-BCS crossover is the superfluid critical temperature, predicted to have a maximum on the BEC side of the strongly interacting regime. However, accurately measuring the critical temperature is challenging due to difficulties in determining a reliable temperature scale in the presence of strong interactions. In this thesis, we determine the critical temperature in the crossover with high accuracy by reconstructing the density distribution and incorporating interaction effects in the low-density wings when fitting to the virial expansion of the equation of state. This requires precise identification of the superfluid phase transition onset, for which we have developed two novel advanced image recognition techniques based on machine learning. Our improved methodology confirms, for the first time, an increase in the critical temperature from the BCS limit, extending beyond the unitarity point and approaching the BEC limit.

Crossing the superfluid phase transition is accompanied by spontaneous symmetry breaking, creating an energy landscape that supports two distinct excitation modes: the Goldstone and Higgs modes. Here, we probe the Higgs mode using two distinct excitation methods: a quench and a modulation of the interaction strength. This enables us to observe the Higgs mode throughout the crossover, revealing a gradual fading of the mode as it approaches the BEC regime, where particle-hole symmetry vanishes. Notably, we observe no temperature dependence of the Higgs mode, prompting further research.

Finally, we present a novel cooling method for a strongly interacting Fermi gas on the BEC side of the crossover, where a composite dimer bound state exists. By applying a modulation of the magnetic field at frequencies close to, but higher than the bound state energy, we selectively dissociate and remove high-energy dimers from the trap, thus realising evaporative cooling of the sample. This method does not require any changes to the trapping potential and facilitates staying in the efficient runaway regime. We demonstrate cooling for a wide range of interactions on the BEC side of the crossover, achieving high efficiencies that match or exceed all previously reported forced evaporation cooling near Feshbach resonances.

Contents

1	Introduction	1
2	Ultracold Fermi gases	9
2.1	Ideal quantum gas	10
2.2	Interatomic interactions	12
2.2.1	Elastic s-wave scattering	13
2.2.2	Feshbach resonances	14
2.2.3	Bound state energy	16
2.3	BEC-BCS crossover	18
2.3.1	BEC regime - Repulsive composite dimers	18
2.3.2	BCS regime - Cooper pairing	19
2.3.3	Superfluid phase transition	25
2.3.4	Collective excitations	26
2.4	Fermions in an inhomogeneous trap	29
3	Apparatus and Preparation	31
3.1	Properties of Lithium and Sodium	32
3.1.1	Transition levels and laser systems	32
3.1.2	Hyperfine splitting of the ground state	33
3.2	Preparation of an ultracold fermionic sample	35
3.2.1	Vacuum system	35
3.2.2	Effusive Oven	35
3.2.3	Zeeman slower	37
3.2.4	Magneto-optical trap	38
3.2.5	Magnetic trap	39
3.2.6	Dipole trap	41
3.2.7	Quantum simulator	43
4	Detection and Methodology	45
4.1	Absorption imaging	46
4.1.1	Imaging system	47
4.1.2	High-intensity absorption imaging	48
4.2	Thermometry of an interacting Fermi gas	52
4.2.1	Inverse Abel transformation - Centre density	53
4.2.2	Reconstruction of the potential	56

4.2.3	Virial expansion of the equation of state - Temperature	57
4.3	Detecting Superfluidity via the rapid ramp technique	59
4.4	Machine Learning	64
4.4.1	Artificial neural networks	65
4.4.2	Supervised learning	69
4.4.3	Unsupervised neural networks	71
4.5	Superimposing a fast magnetic field	72
4.5.1	Quench of the magnetic field	74
4.5.2	Modulation of the magnetic field	74
5	Machine Learning the Superfluid Phase Transition	77
5.1	Detecting the superfluid phase transition	79
5.2	Detecting phase transitions with machine learning techniques	79
5.3	Acquiring datasets across the phase transition	80
5.3.1	Controlled heating of a Fermi gas	80
5.3.2	Datasets	81
5.4	Supervised learning of the superfluid phase transition	84
5.4.1	Training of the neural network	85
5.4.2	Determining the phase transition	90
5.4.3	Critical temperature	90
5.4.4	DeepLift - Insights into the network	91
5.5	Autoencoder - Unsupervised detection of the superfluid phase transition	94
5.5.1	Architecture of the autoencoder network	96
5.5.2	Latent space - Detecting the phase transition in a low-dimensional representation	97
5.5.3	Critical temperature	107
5.5.4	DeepLift in an autoencoder	109
5.6	Conclusion	110
6	Observing the Higgs mode in a strongly interacting Fermi gas	113
6.1	Higgs mode in an inhomogeneous trap	116
6.1.1	Homogeneous trap	116
6.1.2	Inhomogeneous trap	117
6.2	Quenching a superfluid - Time dynamics of the Higgs mode	122
6.2.1	Experimental realisation	124
6.2.2	Observed dynamics	124
6.2.3	Extracted oscillation frequency and damping	128
6.2.4	Conclusion	130
6.3	Parametric excitations of a superfluid - Spectroscopic measurement of the Higgs mode	130
6.3.1	Experimental realisation	131
6.3.2	Observed excitation rate	132
6.3.3	Excitation rate across the crossover	134
6.3.4	Conclusion	137
6.4	Comparison between the excitation techniques	137

7	Cooling of a strongly interacting quantum gas by dimer dissociation	139
7.1	Principle of evaporative cooling	141
7.1.1	Established evaporative cooling techniques	142
7.2	Evaporative cooling via dissociation of composite dimers	144
7.2.1	Evaporative cooling step model	147
7.3	Experimental implementation	148
7.3.1	Observed cooling	150
7.3.2	Onset of dissociation and cooling	153
7.3.3	Frequency of maximum cooling efficiency	154
7.4	Efficiency of cooling	156
7.4.1	Frequency range of efficient cooling	158
7.4.2	Cooling throughout the BEC regime of the BEC-BCS crossover	159
7.4.3	Minimum bound state supporting efficient cooling	160
7.4.4	Efficiency in a gravity-compensated trap	161
7.5	Time scale of the cooling technique	163
7.6	Conclusion	164
8	Summary and Outlook	167
8.1	Summary	167
8.2	Outlook	170
8.2.1	Probing the Higgs mode in a homogeneous Fermi gas	170
8.2.2	The Higgs mode in the crossover from 3D to 2D	171
8.2.3	Measuring the pseudogap	171
A	Box trap setup	175
B	Special functions	179
	Bibliography	181
	Tools	201

Introduction

The discovery of superfluidity

The research on superfluid systems began with the discovery of superconductivity in mercury by Onnes in 1911, for which he was awarded the Nobel Prize in Physics in 1913 [1]. Onnes measured a sharp and sudden drop in electrical resistance at temperatures below a critical temperature of 4.2 K, marking the first observation of superconductivity – an effect that could not be explained by the understanding of physics at the time. Since then, superconductivity has been observed in various other materials [2], including unconventional high temperature superconductors. Today, superconductors are employed in a wide range of applications, from the most powerful magnets [3], e.g., in particle accelerators or MRI devices [4], to high-precision magnetic field probes [5] and superconducting qubits [6].

Superconductivity is directly linked to the broader concept of superfluidity, which describes a system that exhibits zero viscosity [7]. Superconducting systems are essentially superfluid electron gases, allowing an electric current to flow without resistance [8]. The first observation of superfluidity was reported in liquefied helium-4 at temperatures below 2.2 K [9, 10], roughly two decades after superconductivity was discovered in mercury.

Helium-4, being a bosonic species, follows fundamentally different physics compared to a superfluid composed of fermionic electrons in a superconductor. For bosonic systems, a link between Bose-Einstein condensation (BEC) [11] and superfluidity was established, as described by Landau [12]. In a BEC, the lowest energy state becomes macroscopically occupied. If the BEC has finite interactions, the lowest possible excitation is raised to a finite critical velocity, thus enabling superfluidity.

However, this theory does not extend to an ideal Fermi gas, as fermions do not exhibit condensation at low temperatures due to Pauli pressure [8]. Instead, at zero temperature, fermions fill the available energy states up to a maximum energy level known as the Fermi energy, E_F . However, if two interacting Fermi states are introduced, the Pauli exclusion principle no longer prevents them from interacting at low temperatures. This interaction was formalised by Bardeen, Cooper and Schrieffer (BCS) in 1957 [13], who described superfluidity in a two-state mixture of fermions through the condensation of opposite momentum pairs, known as Cooper pairs [14]. In a BCS-type superfluid, the dispersion relation is lifted by an emergent gap parameter Δ , which prohibits excitations with energies or more generally

temperatures below this gap, thus enabling superfluidity.

For Cooper pairs to form, the fermion mixture must support a weak bound state, mediated by attractive interactions. This bound state only exists due to the occupation of fermions up to the Fermi energy E_F , representing a genuine many-body effect [8]. In superconductors, these weak attractive interactions arise from phonon interactions between electrons of opposite spin. Superconductors that follow BCS theory are classified as conventional superconductors.

The BEC-BCS crossover

While BEC- and BCS-type superfluids both conceptually originate from macroscopic condensation, the underlying physics governing bosons and fermions differ significantly. Nevertheless, the shift from a BCS-type to a BEC-type superfluid can be described within a unified framework, known as the BEC-BCS crossover [8, 15, 16]. In the BCS limit, weakly interacting fermions form Cooper pairs, where the pair size is much larger than the average interparticle distance. As the interaction strength increases, the bound state energy of the Cooper pairs grows, reducing the pair size [17] and raising the critical temperature. At a certain interaction strength, the pair size becomes much smaller than the average interparticle distance, allowing the pairs to be treated as composite bosons, which exhibit signatures of BEC. Thus, the shift from the BCS to the BEC limit is not considered a transition but a crossover.

The strongly interacting regime between the BEC and BCS limits is of particular interest. In this regime, the pair size becomes comparable to the interparticle distance. A notable case is the unitarity limit, where the interaction reaches its maximum, and the interparticle distance becomes the only relevant length scale [18]. Interestingly, the highest critical temperature T_C , of all fermionic superfluids is expected to occur in the strongly interacting regime between the BEC and unitarity limits [8, 19, 20]. However, determining the exact position and value of the highest T_C remains a challenge due to the difficulties in modelling the strongly interacting regime.

In theoretical frameworks, simulations of the strongly interacting regime by conventional computational methods are inefficient due to the extensive size of the corresponding Hilbert space [21, 22]. Furthermore, strong interactions preclude the use of perturbative approaches to accurately model the system. As an alternative, condensed matter experiments have been employed to explore the strongly interacting regime, but in such experiments, precise control of the system parameters is highly complex and correlated among the parameters [23].

In this thesis, we leverage the framework of a tailored quantum simulator to address the challenges posed by the BEC-BCS crossover.

Quantum simulation with ultracold gases

Quantum simulators are machines designed to model specific quantum-mechanical systems, enabling the simulation of particular quantum effects in a clean and controlled environment. This precise and clean realisation of quantum systems requires careful disentanglement and meticulous control of environmental parameters, as well as highly accurate readout mechanisms. Ultracold atomic gases offer an ideal platform for simulating quantum-mechanical problems, as they inherently exhibit quantum-mechanical properties.

The emergence of quantum simulators with ultracold atomic gases began with the proposal to cool dilute atomic gases using laser light in 1975 [24]. This led to the development of techniques for cooling and trapping dilute atomic gases via electromagnetic fields, culminating in the Nobel Prize in Physics of 1997 [25–27]. A high degree of control is achieved through the interaction of atoms with electromagnetic fields, which allows for precise tuning of the density, atom number, and temperature of the samples. Furthermore, the interactions between different atomic states can be controlled by varying an external magnetic field, a phenomenon known as Feshbach resonances [28, 29].

The implementation of these techniques led to the observation of the first BEC in bosonic ^{87}Rb [30], and the creation of a degenerate Fermi gas in ^{40}K [31]. By utilising the two-body bound state in a Feshbach resonance, bosonic dimers from two-state mixtures of fermions were created [32, 33] and observed to undergo a phase transition to a BEC [34–36]. Shortly afterward, strongly correlated states such as the Mott insulator phase [37] and superfluidity in the unitary Fermi gas [38, 39] were observed. This marks the beginning of research into the BEC-BCS crossover using quantum simulators.

Our apparatus utilises a mixture of fermionic ^6Li and bosonic ^{23}Na . Both species are first simultaneously cooled down and trapped in a magneto-optical trap [40, 41]. After the fermionic ^6Li is sympathetically cooled through the complete evaporation of the bosonic ^{23}Na , the degenerate fermionic sample is loaded into the final optical dipole trap. Here we create a two-state (hyperfine) mixture of ^6Li , which is manipulated and probed using a combination of carefully designed electromagnetic fields, while interactions are tuned using a Feshbach resonance. This setup provides a platform to explore superfluidity within the BEC-BCS crossover.

Superfluid critical temperature in the BEC-BCS crossover

The most prominent feature of the BEC-BCS crossover is the onset of superfluidity. As discussed earlier, determining the superfluid critical temperature in the strongly interacting regime between the BEC and BCS limits is remarkably non-trivial, and an exact determination has thus far been elusive.

In the BEC limit, the superfluid critical temperature corresponds to the critical temperature of the BEC transition, with a small correction to higher critical temperatures for stronger interactions [42, 43]. Conversely, in the BCS limit, the critical temperature decreases exponentially as the interaction weakens [44]. The critical temperatures in both limits is sketched in Fig. 1.1. Connecting both limits in the sketch suggests a maximum of the critical temperature near the unitarity regime. However, no exact theory for the critical temperature in the strongly interacting regime exists, and both the existence and location of a maximum in the critical temperature remain topics of ongoing debate [19, 20, 45–47].

In this thesis, we measure the phase diagram of the superfluid critical temperature in the strongly interacting regime of the BEC-BCS crossover with unprecedented accuracy. Accurately determining the temperature at the onset of superfluidity requires a precise detection and measurement of the onset of condensation. In the BEC limit, condensation appears as a sharp bimodal peak atop a thermal background during imaging after a brief expansion time [48]. However, this feature is absent in the BCS limit, where weakly bound

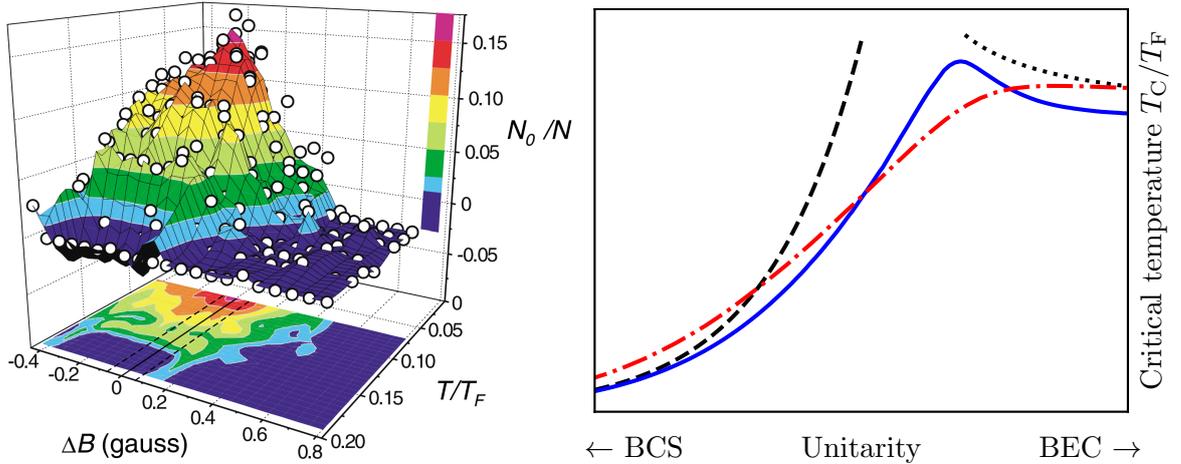


Figure 1.1: Superfluidity in the BEC-BCS crossover. The figure depicts condensation and superfluidity in the BEC-BCS crossover. **Left:** An early measurement of condensation in the crossover regime, adapted from [38]. **Right:** The critical temperature T_C/T_F of superfluidity in the crossover regime, connecting the critical temperatures in the BEC (dotted line) [42, 43] and BCS (dashed line) limits [44]. Around the unitarity regime of strong interactions, no exact solution exists. Several theoretical approaches are available, e.g., the red dashed-dotted line [45] and the blue solid line [19]. More details about the phase diagram are given in Fig. 2.7.

Cooper pairs break during expansion. To address this, early measurements throughout the crossover employed the rapid ramp technique [38, 49], which projects weakly bound Cooper pairs onto tightly-bound dimers during expansion, restoring a bimodal momentum distribution. This established the first phase diagram of superfluidity in the BEC-BCS crossover, shown on the left side of Fig. 1.1. However, the rapid ramp method introduces its own set of complexities and challenges [50–52].

In the past decade, neural networks have garnered increasing attention, being integrated into a wide range of technologies, from large language models [53] to image generation via diffusion models [54, 55]. Over this time, the application of machine learning techniques in physics has also proven advantageous [56, 57], which is further indicated by the recent Nobel Prize in Physics of 2024 [58]. For instance, machine learning has been employed in optimisation tasks [59, 60] and in detecting phase transitions in both theoretical [61–64] and experimental data [65, 66]. This thesis presents the first experimental detection of superfluidity in the entire BEC-BCS crossover, without employing conventional methods, such as the rapid ramp technique.

To precisely determine the critical temperature of condensation, we have developed an advanced image recognition technique that directly infers condensation from standard time-of-flight images. This technique utilises image recognition via deep convolutional neural networks. We have developed two distinct neural network architectures. The first architecture relies on manually labelled data from the rapid ramp technique, enabling the network to predict condensate fractions from previously unseen time-of-flight images after training [67]. The second architecture is a deep convolutional autoencoder neural network that requires neither labelling nor the rapid ramp [68]. After carefully determining the thermometric

parameters at the onset of superfluidity – such as temperature and density – this approach has allowed us to map the superfluid phase diagram in the strongly interacting regime, experimentally confirming, for the first time, a maximum critical temperature on the bosonic side of the crossover [67, 68].

Superfluids out of equilibrium

When crossing the superfluid phase transition from an initially thermal gas, the system’s order parameter Δ acquires a finite value, representing the coherent macroscopic wavefunction of the condensate [69]. The free energy landscape of a system as a function of the order parameter Δ is described by the Ginzburg-Landau theory [70]. This energy landscape resembles a “Mexican hat” and features a minimum energy along a ring of constant $|\Delta|$, independent of the phase of Δ , which describes a continuous symmetry. As the order parameter takes on a finite value, it arbitrarily selects a phase, thereby breaking the symmetry.

Such systems exhibit two fundamental excitation modes. The first is the (Nambu-)Goldstone mode, which alters the phase of the order parameter along the ring-like valley of minimum energy [71, 72]. This mode manifests as phonons and is well understood. More intriguing is the second excitation mode, known as the Higgs mode [70], which modulates the absolute value of the order parameter Δ and requires a minimum energy of 2Δ to be excited [73].

Exciting the Higgs mode is non-trivial because there is no direct control over the order parameter Δ . Instead, excitation is achieved by driving the system out of equilibrium, creating a mismatch between Δ and the system’s equilibrium gap parameter. This allows for two distinct excitation methods: sudden quenches [73–77] and parametric excitations [78–83] of the order parameter.

The Higgs mode reveals itself as an oscillation in the order parameter with a frequency of $\omega_H = 2\Delta/\hbar$ [73]. However, measurements of the Higgs mode in the strongly interacting regime are extremely challenging, and only a few such measurements have been reported. The earliest measurement in this regime was conducted by my predecessors in our group [78]. Recently, a measurement of the Higgs mode at unitarity has been published [75]. A precise understanding of the Higgs mode throughout the crossover requires more precise measurements.

In an interacting Fermi gas, the Higgs mode is predicted to be stable against decay into lower-lying modes only in the BCS limit [84], where stability is mediated by particle-hole symmetry, giving the system an effective Lorentz invariance [71]. This particle-hole symmetry is a feature of the BCS limit and gradually diminishes as one approaches the BEC limit. Additionally, the Higgs mode is expected to show a strong temperature dependence due to the diminishing condensate as the system approaches the critical temperature T_C . However, the recent measurement of the Higgs mode at unitarity observed no significant temperature dependence [75].

In this thesis, we probe the Higgs mode using both of the excitation methods discussed earlier. Following a sudden quench of the order parameter, we observe time-resolved oscillation of the order parameter throughout the BEC-BCS crossover, for the first time. By approaching the BEC limit, the oscillating signature fades out, signalling an increasing instability of the Higgs mode due to diminishing particle-hole symmetry. Additionally, the Higgs mode centre

frequency and width are probed by parametric excitations. While this confirms an averaging of the Higgs mode centre frequency in an inhomogeneous trapping potential, we observe no temperature dependence of the centre frequency. This confirms the recent measurement at unitarity [75], yet cannot be explained by current theories.

Cooling by dimer dissociation

As discussed earlier, the advent of experiments on quantum simulators using ultracold atomic gases began with the development of laser-based cooling and trapping techniques [24]. To reach the ultra-low temperatures required for observing strongly correlated systems, such as the Mott insulator phase [37] or condensation in the unitary Fermi gas [38, 39], evaporative cooling was introduced. Evaporative cooling selectively removes the highest energy particles from a sample, thereby reducing the average energy of the remaining particles and effectively lowering the temperature [85]. Although this process decreases both the temperature and the number of atoms, efficient evaporation increases the phase-space density of the sample, thus lowering its the entropy.

Generally speaking, evaporative cooling can be implemented in two different ways: in magnetic traps and in optical dipole traps. In magnetic traps, atoms are selectively evaporated by driving radio-frequency-transitions (RF) to high-field-seeking states, which are anti-trapped and expelled from the system [86]. However, magnetic traps can only confine specific atomic states, making them unsuitable for most quantum simulators involving Fermi gas state mixtures, where more flexible state preparation is required. Optical dipole traps address this limitation, though they lack a selective mechanism for removing high-energy atoms from the trap [87]. Instead, the optical dipole trap is ramped down in a controlled manner, which primarily releases the highest energy atoms from the trap but also reduces the overall density of the system, slowing down thermalisation. Although, both methods are widely used in ultracold atom experiments, they have inherent limitations.

Theoretically, even more complex and interesting quantum states are predicted to emerge at lower temperatures, which are not yet achievable with current cooling methods. These states include unconventional superconductivity [88], topological quantum states [89], and spin liquids [90]. Therefore, the development of novel cooling techniques remains crucial for enabling the observation of these peculiar quantum phases.

In this thesis, we introduce a novel cooling method based on the selective dissociation of high-energy dimers. On the BEC side of the Feshbach resonance, bosonic dimers exist with a bound state energy, E_B . These dimers can be dissociated through parametric excitation at frequencies higher than their corresponding bound state energy. Dissociating a dimer produces two unpaired fermions, which interact more strongly with the remaining dimer cloud. This interaction increases the repulsive mean-field energy, thereby raising the threshold dissociation energy beyond E_B , by an amount proportional to the dimer density. Consequently, high-energy dimers near the dilute edge of a harmonic trap require less energy to dissociate compared to low-energy dimers in the dense central region. Once a dimer is broken, the resulting fermions gain excess kinetic energy, which may expel them from the trap. This provides a mechanism to selectively remove high-energy dimers from the system, cooling the sample.

We explore this cooling mechanism throughout the BEC side of the crossover, from unitarity to deep bound state energies of > 200 kHz. Our method achieves high evaporation efficiencies γ , with values of $\gamma = 4$ in a gravity-tilted trap and $\gamma = 6$ in a gravity-compensated trap, matching and even surpassing the highest reported efficiencies for forced evaporation in optical dipole traps near a Feshbach resonance [85, 87, 91–99].

Outline of this thesis

The main results of this thesis are the detection of the superfluid phase transition via unsupervised machine learning, detailed in Chapter 5, and the development and demonstration of a novel cooling technique through dimer dissociation, as described in Chapter 7. To provide context and elaborate on the underlying physics and analysis techniques, the thesis is structured as follows:

- Chapter 2: This chapter lays out the theoretical framework, presenting the essential concepts and ideas necessary for understanding the content of this thesis.
- Chapter 3: Here, we focus on the experimental apparatus and the preparation of the degenerate Fermi spin mixture that constitutes our quantum simulator.
- Chapter 4: This chapter outlines the detection techniques and methodologies used throughout the thesis. It covers imaging, excitations via a superimposed magnetic field, thermometry, and the detection of superfluidity via both conventional methods and machine learning techniques.
- Chapter 5: In this chapter, we present the first key result of this thesis: the superfluid phase diagram, mapped using novel image recognition techniques based on machine learning.
- Chapter 6: This chapter is dedicated to the observation of the Higgs mode in the strongly interacting regime of the BEC-BCS crossover.
- Chapter 7: We introduce the newly developed cooling method based on dimer dissociation in this chapter, which presents the second main result of this thesis.
- Chapter 8: The final chapter summarises the results and offers a brief outlook on future directions.

Publications

Parts of this thesis have been or will be published in the following articles:

- M. Link, K. Gao, A. Kell, M. Breyer, **D. Eberz**, B. Rauf and M. Köhl – *Machine Learning the Phase Diagram of a Strongly Interacting Fermi Gas* – Published in Physical Review Letters [67] – DOI: 10.1103/PhysRevLett.130.203401

- **D. Eberz**, M. Link, A. Kell, M. Breyer, K. Gao and M. Köhl – *Detecting the phase transition in a strongly interacting Fermi gas by unsupervised machine learning* – Published in Physical Review A [68] – DOI: 10.1103/PhysRevA.108.063303
- A. Kell, M. Breyer, **D. Eberz** and M. Köhl – *Exciting the Higgs Mode in a Strongly Interacting Fermi Gas by Interaction Modulation* – Published in Physical Review Letters [100] – DOI: 10.1103/PhysRevLett.133.150403
- M. Breyer, **D. Eberz**, A. Kell and M. Köhl – *Quenching a Fermi superfluid across the BEC-BCS crossover* – Under review at SciPost Physics – Preprint available in [101]
- **D. Eberz**, A. Kell, M. Breyer and M. Köhl – *Cooling a strongly-interacting quantum gas by interaction modulation* – Under review at Physical Review Letters – Preprint available in [102] – DOI: 10.48550/arXiv.2410.10642

Ultracold Fermi gases

This chapter provides an introduction to the concepts and theories essential for understanding this thesis. The experiments presented in this thesis are based on a quantum simulator using ultracold ${}^6\text{Li}$, a fermionic isotope of lithium. Thus, the first section introduces the fundamental concepts of an ideal Fermi gas, as described by the Fermi-Dirac distribution. This will be compared to the intrinsic statistics of an ideal Bose gas and the classical description provided by the Boltzmann distribution.

As interactions freeze out in a polarised ultracold Fermi gas, we introduce a two-state mixture and discuss interatomic interactions between distinguishable fermions. Scattering between distinguishable fermions introduces the concept of (Fano-)Feshbach resonances, which serve as a “tuning knob” to adjust interatomic interactions in the experiment by varying the external magnetic field – a crucial feature for realising various quantum states.

By adjusting the interaction strength, the Fermi gas can be tuned across the crossover from tightly-bound dimer states to the many-body regime of Cooper pairing on the Fermi surface, known as the BEC-BCS crossover. To understand the physics within this crossover, we will briefly review its limits and illustrate the smooth evolution of pairing from Cooper pairs to tightly-bound composite dimers. Here, we discuss key aspects of the crossover, including condensation, superfluidity, and the excitation spectrum.

In the experiment, the atoms are prepared in an inhomogeneous trapping potential. Thus, the final section provides a brief overview of fermions in external potentials.

For more comprehensive reviews of the concepts covered in this chapter, refer to various review papers and books such as [8, 18, 48, 103, 104].

Outline of this chapter

The sections of this chapter are outlined as follows:

- Section 2.1: We begin by introducing the intrinsic quantum statistics of ideal Bose and Fermi gases, highlighting the differences to a classical gas described by Boltzmann statistics.
- Section 2.2: Non-ideal quantum gases exhibit interesting effects when interactions are introduced. This section covers the principles of interactions and introduces the concept

of Feshbach resonances – a key mechanism to tune interactions freely. Interactions are also linked to the existence of a two-body bound state, which we discuss further.

- Section 2.3: By tuning interactions in a two-state Fermi gas, the system can be tuned from a molecular Bose-Einstein condensate (BEC) to a condensate of Cooper pairs (BCS). This section describes the progression of physics from the BEC to the BCS regime, known as the BEC-BCS crossover, and covers the effects of pairing on condensation, the critical temperature and the emergence of collective modes, such as the Higgs mode, across the crossover.
- Section 2.4: In our system, the gas is confined within an inhomogeneous optical dipole trap potential, approximated by a 3D harmonic potential. Here, we introduce the local density approximation (LDA), which describes an inhomogeneous sample through integrals over homogeneous slices.

2.1 Ideal quantum gas

At high temperatures T , particles in a gas behave classically, with their energies E governed by the Boltzmann distribution [8]

$$f_{\text{Boltzmann}}(E, T) \propto \exp\left(-\frac{E}{k_{\text{B}}T}\right), \quad (2.1)$$

where k_{B} is the Boltzmann constant. As the temperature approaches the quantum regime, the particles behaviour diverges based on their intrinsic spin. Particles with half-integer spin are classified as “Fermions”, while those with integer spin are called “Bosons”. According to the spin-statistics theorem, the total wave function $\psi(r_1, r_2)$ of multiple indistinguishable particles remains unchanged under exchange of bosons (symmetric), while it collects a factor of -1 under the exchange of indistinguishable fermions (antisymmetric) [105].

This small distinction has significant implications for indistinguishable fermionic particles: if they occupy the same position $r_1 = r_2$, the only possible solution for the wave function is $\psi(r_1, r_2) = 0$. Thus, two indistinguishable fermionic particles can never occupy the same location (or, more broadly, the same state), a principle known as the “Pauli exclusion principle”, which underlies the behaviour of our fermionic ${}^6\text{Li}$ sample.

The Pauli exclusion principle gives rise to the Fermi-Dirac distribution [8]

$$f_{\text{Fermi}}(E, T) = \frac{1}{\exp\left(\frac{E-\mu}{k_{\text{B}}T}\right) + 1}, \quad (2.2)$$

which describes the occupation probability of a state E at temperature T with a chemical potential $\mu(T)$. At zero temperature $T = 0$, all states up to $\mu(T = 0)$ are filled, while states at higher energies remain unoccupied, a phenomenon also known as the Fermi edge. Here, the chemical potential equals the Fermi energy, $\mu(T = 0) = E_{\text{F}}$.

As temperature rises, the distribution loses its step-function-like behaviour and softens near the Fermi edge, with some particles occupying higher energy states and some states below the Fermi becoming vacant. Consequently, as new particles can now occupy lower-lying

energy states $E < E_F$, the chemical potential $\mu(T)$ decreases with increasing temperature T . The temperature dependence of the chemical potential $\mu(T)$ is self-consistently determined by fixing the atom number

$$N = \int dE \rho(E) f(E, T), \quad (2.3)$$

where $\rho(E)$ is the density of states of the corresponding trapping potential. In Fig. 2.1, the Fermi-Dirac distribution from Eq. (2.2) (left) and the self-consistently calculated chemical potential $\mu(T)$ (right) are shown for various temperatures in a homogeneous gas ($\rho(E) \propto \sqrt{E}$), with values in terms of the Fermi energy E_F and Fermi temperature $T_F = E_F/k_B$. As the temperature approaches the order of the Fermi temperature T_F , the Fermi-Dirac distribution noticeably “melts”, converging towards the Boltzmann distribution (dashed line, left plot).

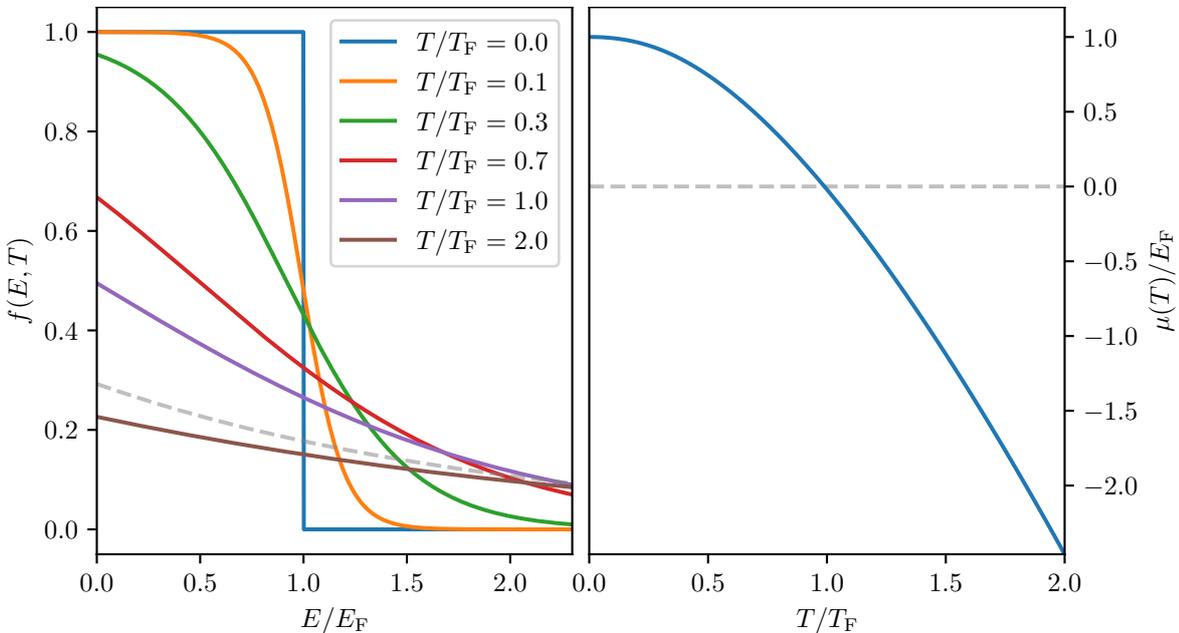


Figure 2.1: Fermi statistics **Left:** The Fermi-Dirac distribution from Eq. (2.2) is shown for increasing temperatures T (solid lines) in a homogeneous system. As temperature increases, the distribution approaches the Boltzmann distribution (grey dashed line), shown for the highest temperature of $T/T_F = 2.0$. **Right:** Self-consistently determined temperature dependence of the chemical potential $\mu(T)$ in a homogeneous Fermi gas, calculated by fixing the atom number as in Eq. (2.3). All temperatures and energies are in units of the Fermi energy and temperature (E_F , T_F).

In a homogeneous gas, where $E = (\hbar k)^2/(2m)$, the Fermi energy E_F can directly be related to the density per spin state n_σ by integrating the Fermi-Dirac distribution (Eq. (2.2)) at

$T = 0$ over the momentum space $d^3\mathbf{p}$, resulting in [8]

$$\begin{aligned}
 n_\sigma &= \int \frac{d^3\mathbf{p}}{(2\pi\hbar)^3} f_{\text{Fermi}}\left(\frac{p^2}{2m}, T = 0\right) = 4\pi \int_0^{\sqrt{2mE_F}} \frac{dp}{(2\pi\hbar)^3} p^2 = \frac{1}{6\pi^2} \left(\frac{2mE_F}{\hbar^2}\right)^{\frac{3}{2}} \\
 &\Leftrightarrow E_F = \frac{\hbar^2}{2m} (6\pi^2 n_\sigma)^{\frac{2}{3}} \\
 &\Leftrightarrow k_F = (6\pi^2 n_\sigma)^{\frac{1}{3}},
 \end{aligned} \tag{2.4}$$

where, in the final step, we introduce the maximal occupied momentum state, namely the Fermi momentum k_F , with m denoting the particles' mass. Note that in order to calculate the density profile in an inhomogeneous trap with potential $V(\mathbf{r})$, the chemical potential is replaced by a spatially varying chemical potential $\mu \rightarrow \mu - V(\mathbf{r})$, a technique known as the ‘‘local density approximation’’, further explored in Section 2.4.

While we conduct our experiments on fermionic ${}^6\text{Li}$, the experimental apparatus also hosts bosonic ${}^{23}\text{Na}$, which is used to sympathetically cool lithium (refer to Section 3.2.5). Furthermore, as we will see in Section 2.2.3, the fermionic particles can also form a bosonic two-body bound state. For bosons, the corresponding distribution function is the Bose-Einstein distribution, which differs only by a change of sign in the denominator [8]

$$f_{\text{Bose}}(E, T) = \frac{1}{\exp\left(\frac{E-\mu}{k_B T}\right) - 1}, \tag{2.5}$$

where the chemical potential must be smaller than the lowest available energy state $\mu \leq E_0$ (ground state), to ensure $f_{\text{Bosons}} \geq 0$. Due to the sign change, the distribution displays divergent behaviour at zero temperature $T = 0$ or $E = \mu$, allowing for a macroscopic occupation of the ground state – a phenomenon known as ‘‘Bose-Einstein condensation’’.

Bose-Einstein condensation occurs when all accessible excited states with $E > E_0$ are occupied and all remaining atoms $N_0 = N - N_{\text{exc}}$ enter the ground state. As the number of available excited states N_{exc} decreases along with temperature, there exists a critical temperature T_C , marking the onset of the condensation.

2.2 Interatomic interactions

To observe rich and complex fermionic phenomena, interactions play a crucial role. As the Pauli exclusion principle prevents interactions in a polarised sample of ultracold fermions, we use a spin mixture. In our experiment, this mixture consist of two of the three lowest magnetic hyperfine states of the ground state $2^2\text{S}_{1/2}, F = 1/2$ in ${}^6\text{Li}$.

The scattering process between these states is governed by the central interatomic potential $V(r)$. At large interatomic distances, the potential asymptotically approaches the attractive, long-range Van-der-Waals interaction ($-C_6/r^6$) [106]. At very short distances, the Coulomb repulsion between the electron clouds forms a hard-core potential wall, further reinforced by Pauli pressure for aligned spins in a triplet state [106]. In the intermediate range, the exchange integral for spin singlet and triplet states influences the potential [103]. While the exact solution of the two-body problem via the potential would be extremely difficult, the

ultracold gases in the experiment have extremely large intrinsic length scales, simplifying the physics enormously.

Since the potential asymptotically connects to the Van-der-Waals potential, which has a natural length scale $(2R_{\text{vdW}} = (2\mu_{\text{r}}C_6/\hbar^2)^{1/4})$ [106], we can compare the extent of the interatomic potential to the mean interatomic distance. For ${}^6\text{Li}$, the Van-der-Waals coefficient is determined to be $C_6 = 1393.39(16)E_h a_0^6$ [107], where a_0 is the Bohr radius and E_h is the Hartree energy. Here, the reduced mass $\mu_{\text{r}} = m_{\text{Li}}/2$ is half the mass of a lithium atom, m_{Li} . This yields a characteristic length scale of $R_{\text{vdW}} = 31.26 a_0 \approx 1.7 \text{ nm}$.

In contrast, the mean interatomic distance for dilute ultracold samples is determined by the inverse density $1/n^{1/3} \propto k_{\text{F}}^{-1} \approx 10^4 a_0$ [8], estimated for typical densities of $n \approx 10 \mu\text{m}^{-3}$. Additionally, the experimental temperatures are well below the Fermi energy $T/T_{\text{F}} \ll 1$, placing the de-Broglie wavelength on the same order of magnitude [8, 108].

Since the interatomic distance is significantly larger than the range of the potential $k_{\text{F}}^{-1} \gg R_{\text{vdW}}$, the fine details of the scattering dynamics become unresolved and can thus be encapsulated in a single parameter: the scattering length a . The scattering length a will be introduced quantitatively in the next section.

2.2.1 Elastic s-wave scattering

To solve the scattering problem between two particles, the common approach is to switch to centre-of-mass coordinates. In this frame, the problem is expressed in terms of the reduced mass $\mu_{\text{r}} = m_{\text{Li}}/2$ and the relative coordinates $\mathbf{r} = \mathbf{r}_1 - \mathbf{r}_2$ of the two particles ($\mathbf{r}_{1,2}$), with the relative distance defined as $r := |\mathbf{r}|$. The Schrödinger equation for the interatomic potential $V(\mathbf{r})$ becomes [8, 106, 108]

$$\left(\frac{\hbar^2}{2\mu_{\text{r}}} \nabla_{\mathbf{r}}^2 + V(\mathbf{r}) \right) \psi_{\mathbf{k}}(\mathbf{r}) = \frac{\hbar^2 k^2}{2\mu_{\text{r}}} \psi_{\mathbf{k}}(\mathbf{r}), \quad (2.6)$$

where $k = \sqrt{2\mu_{\text{r}}E}/\hbar$ represents the wave vector associated with the energy E of the incident wave. Since we consider only elastic scattering among the states involved, the wave vector of the outgoing wave \mathbf{k}' must have the same energy. Moreover, in the ultracold regime, where all wave vectors are much smaller than the inverse of the range of the interatomic potential, only s -wave scattering is relevant.

This allows for an ansatz in the asymptotic regime ($r \gg R_{\text{vdW}}$) of the form [8, 103, 106]

$$\psi_{\mathbf{k}}(\mathbf{r}) \propto \exp(i\mathbf{k} \cdot \mathbf{r}) + f_s(k) \frac{\exp(i\mathbf{k} \cdot \mathbf{r})}{r}, \quad (2.7)$$

where the first term represents the incident plane wave and the second term is the scattered spherical wave with s -wave amplitude $f_s(k)$. The amplitude of the s -wave scattered spherical wave can be written as [8, 103, 106]

$$f_s(k) = \frac{1}{-\frac{1}{a} + r_{\text{eff}} \frac{k^2}{2} - ik}, \quad (2.8)$$

where r_{eff} denotes the effective range, typically on the order of the Van-der-Waals length

scale R_{vdW} , and a being the scattering length.

In the limit of large interatomic distances, where $k|a| \ll 1$ and $r_{\text{eff}} \lesssim 1/k$, the scattering amplitude converges to the scattering length itself, $f_s \rightarrow -a$. Here, the scattering cross-section σ approaches that of scattering from hard cores of radius a , giving $\sigma = 4\pi a^2$.

In contrast, when the scattering length is very large compared to the interatomic distance $k|a| \gg 1$, the scattering amplitude converges to $f_s \rightarrow i/k$, resulting in a scattering cross-section of $\sigma = 4\pi/k^2$. In this “unitarity” limit, the scattering length a – and therefore any details about the interaction potential – become irrelevant and drop out, leaving the interatomic distance as the only pertinent length scale. The effective range, serving as a correction term to this universal behaviour, can typically be neglected for $r_{\text{eff}} \ll 1/k$, as discussed in the previous section.

2.2.2 Feshbach resonances

The phenomenon of (Fano-)Feshbach resonances [28, 109] describes the tunability of the two-state scattering length a via an external magnetic field. This concept is illustrated in Fig. 2.2 and will be discussed in this section. For more detailed discussions, consult [8, 103, 106, 110].

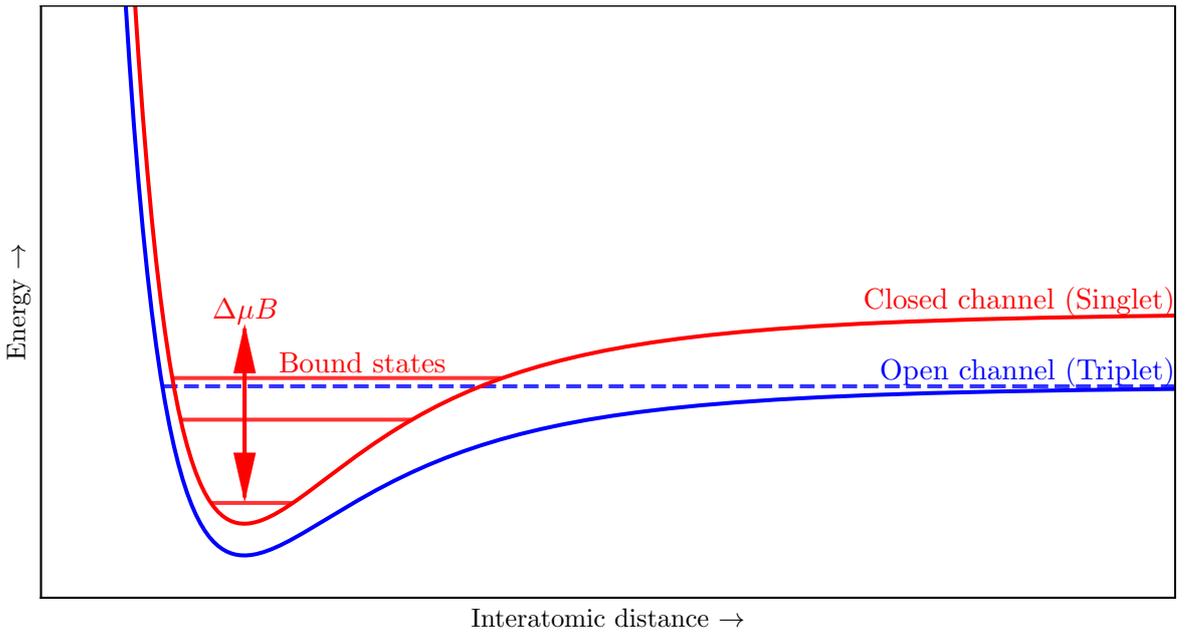


Figure 2.2: Concept of Feshbach resonances. A spin mixture is prepared in the open (triplet) channel (blue). By applying a magnetic field B , a bound state of the inaccessible closed (singlet) channel (red) is tuned into resonance with the incoming energy of the free atoms (dashed line). The shift is enabled by a difference in the channels’ magnetic moment $\Delta\mu$. Sketch inspired by [110].

Normally, the scattering length a is fully described by the interatomic potential $V(r)$, giving rise to a background scattering length a_{bg} [106], which depends on the prepared state mixture. This prepared state mixture is known as the “open channel”.

In the case of ${}^6\text{Li}$, the hyperfine splitting of the ground state ($2^2\text{S}_{1/2}$, $L = 0 \rightarrow m_J = m_S$) gives rise to a high-field-seeking manifold $|1\rangle - |3\rangle$ ($F = 1/2$, $m_S = -1/2$), and a low-field-seeking manifold $|4\rangle - |6\rangle$ ($F = 3/2$, $m_S = 1/2$). A detailed illustration of the level structure is provided in Fig. 3.2. In our experiment, the atomic sample is prepared in a two-spin mixture within the high-field-seeking manifold, meaning the open channel is in a triplet state with a total spin projection of $m_{S,1} + m_{S,2} = -1$.

While the hyperfine splitting of the corresponding ground state allows for several outgoing channels, all allowed scattering channels are “closed channels” because their asymptotic energies, defined by the Zeeman splitting, lie far above the incoming energy. However, due to the differences in spin alignments, the states of the closed channels have distinct magnetic moments μ , leading to a difference of $\Delta\mu$. This difference enables the relative energy of these channels to be adjusted by $\Delta\mu B$, which is realised by changing the external magnetic field B .

Additionally, the spin-nuclear coupling term $\hat{\mathbf{S}} \cdot \hat{\mathbf{I}}$ introduces mixing between singlet and triplet states. Now, if the external magnetic field B shifts a bound state of a previously closed channel close to the incoming energy of the open channel as illustrated in Fig. 2.2, strong coupling occurs, significantly altering the scattering length a . If the bound state energy is tuned slightly above the energy of the incoming channel, the scattering length a becomes large and negative; if it is tuned below, a becomes large and positive. On resonance, when the bound state energy matches exactly, the scattering length diverges, realising the “unitarity regime”, previously discussed in Section 2.2.1.

The behaviour of the scattering length a near resonance can be parametrised in the usual form[106, 111]

$$a(B) = a_{\text{bg}} \left(\frac{\Delta B}{B - B_0} \right), \quad (2.9)$$

where B is the applied magnetic field, ΔB the resonance width and B_0 the resonance position in terms of the magnetic field.

Since interatomic potentials are generally not fully known, Feshbach resonances are determined experimentally. The most recent and precise measurements for resonances in ${}^6\text{Li}$ are provided by [111] and are listed in Table 2.1 in terms of the given parametrisation from Eq. (2.9). For accuracy in describing the behaviour far from resonance, such as zero crossings, direct measurements of the scattering length are preferred. Accordingly, [111] provides measured values of the scattering length a for external magnetic fields B from 1 to 1000 G for all spin mixtures among the $|1\rangle - |3\rangle$ states, plotted in Fig. 2.3 for our experimentally relevant magnetic fields. For further analysis in this thesis, these direct measurements are used.

Table 2.1: Feshbach resonance values from [111]. The table shows the experimentally measured values of the three Feshbach resonances among the $|1\rangle - |3\rangle$ states. The parameters match the parametrisation in Eq. (2.9). a_0 denotes the Bohr radius.

	$ 12\rangle$	$ 13\rangle$	$ 23\rangle$
B_0/G	832.18(8)	689.68(8)	809.76(5)
$\Delta B/\text{G}$	262.3(3)	166.6(3)	200.2(5)
a_{bg}/a_0	-1582(1)	-1770(5)	-1642(5)

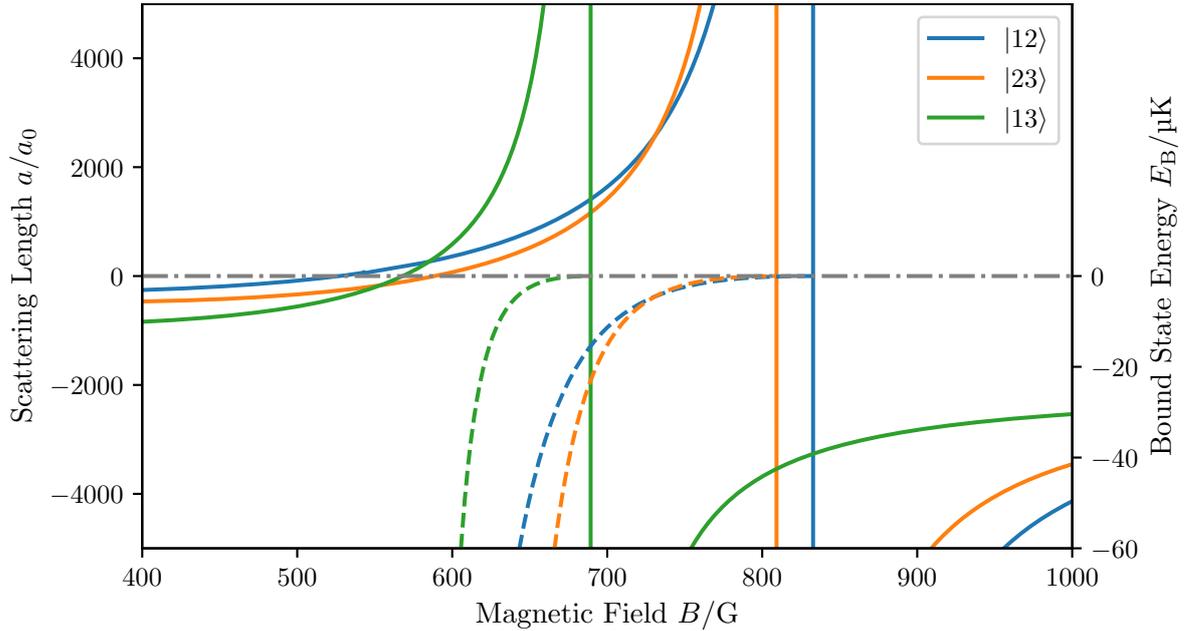


Figure 2.3: Feshbach resonances in ${}^6\text{Li}$. The figure shows the scattering length a/a_0 (solid line) and the corresponding bound state energy E_B (dashed line) as functions of the magnetic field B/G . Here, a_0 denotes the Bohr radius. Data for the scattering length a is taken from [111]. The bound state energy is calculated via Eq. (2.11).

2.2.3 Bound state energy

A key feature of the diverging scattering length at resonance is the emergence of a molecular (dimer) bound state. In the case of Feshbach resonances, this occurs when the bound state energy of the “closed channel” aligns with the energy of the incoming “open channel”. Due to the introduction of coupling, the new dressed states display an avoided crossing, featuring a dimer bound state. Here, the bound state energy asymptotically connects to the energy (relative to the “open channel” energy) of the bound state in the “closed channel” [103, 106].

The bound state appears at unitarity $B = B_0$ and exists for positive scattering lengths $a > 0 \Rightarrow B < B_0$. In the limit of a very large scattering length $k|a| \gg 1$, the bound state energy can be expressed as (often referred to as the effective-range result) [106]

$$E_B = -\frac{\hbar^2}{2\mu_r a^2}, \quad (2.10)$$

where $\mu_r = m_{\text{Li}}/2$ is the reduced mass of the two-body scattering and a is the scattering length. Since the scattering length is usually much larger than the potential’s effective range $a \gg r_{\text{eff}}$, the bound state indicates that it occupies a very high rovibrational molecular state. Nevertheless, these states exhibit very long lifetimes [32, 92, 112] as three-body induced relaxation into deeper lying states is strongly suppressed by Pauli blocking, unlike Feshbach molecules in bosonic species where significant inelastic losses have been observed [29, 113].

While the bound state energy in Eq. (2.10) provides a good approximation near resonance, it lacks details of the scattering potential such as the effective range r_{eff} . For a more precise picture, we consider the Van-der-Waals potential ($-C_6/r^6$) in the asymptotic limits, as discussed in Section 2.2. Considering this limit, a more careful analysis of the scattering potential has been conducted in [114], leading to a higher-order correction of the bound state energy in [115]. The higher-order correction is given up to second order as [115]

$$E_{\text{B}} = -\frac{\hbar^2}{2\mu_{\text{r}}(a-\bar{a})^2} \left[1 + \frac{g_1\bar{a}}{a-\bar{a}} + \frac{g_2\bar{a}^2}{(a-\bar{a})^2} + \mathcal{O}\left(\frac{\bar{a}^3}{(a-\bar{a})^3}\right) \right]. \quad (2.11)$$

In this correction μ_{r} is the reduced mass, a the scattering length, \bar{a} the mean scattering length and $g_{1,2}$ correction factors. The mean scattering length \bar{a} presents an alternative convention of the effective range introduced in [116]. For the Van-der-Waals range R_{vdW} , the mean scattering length is $\bar{a} = \cos(\frac{\pi}{4}) \frac{\Gamma(3/4)}{\Gamma(5/4)} \cdot R_{\text{vdW}} \approx 29.9 a_0$, where a_0 is the Bohr radius. The correction factors are provided in [115] as $g_1 = 0.9179195$ and $g_2 = -0.9467798$.

Both the effective-range result in Eq. (2.10) and the higher-order correction result in Eq. (2.11) are presented in Fig. 2.4 as functions of the scattering length a .

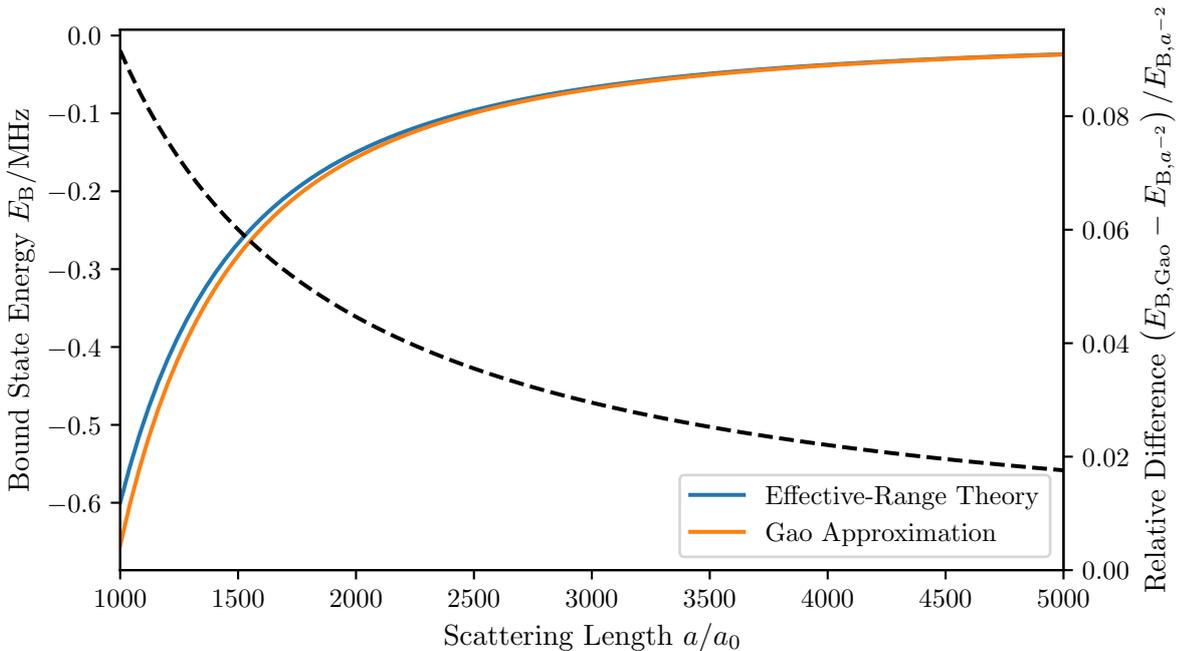


Figure 2.4: Feshbach bound state energy. The figure shows the bound state energy of a Feshbach resonance in terms of the scattering length a/a_0 for the simple effective-range result in Eq. (2.10) and the higher-order correction in Eq. (2.11) derived by [115] (bold lines). Additionally, the relative difference is shown (dashed line).

While both results (solid lines) converge towards each other approaching unitarity ($a \rightarrow \infty$), the relative difference (dashed line) reveals that the simple effective-range result underestimates the bound state energy. In the experiment, we employ scattering lengths as small as $a/a_0 \approx 1500$, where the corrected result gives up to 6% deeper bound state energies. The

corrected result will become relevant later in Chapter 7.

2.3 BEC-BCS crossover

In this section, we discuss the effect of varying the scattering length a across resonance in a two-state mixture of a Fermi gas. The scattering length a can easily be tuned in the experiment using Feshbach resonances, as introduced in Section 2.2.2. This enables us to explore a plethora of interactions and pairing mechanisms by simply tuning the magnetic field.

As introduced in Section 2.2.3, Feshbach resonances introduce a two-particle bound state on the side of positive scattering lengths $a > 0$, forming composite dimers. If the bound state energy of these composite dimers is much larger than all other energy scales of the gas, these dimers can be considered stable entities, effectively obeying Bose statistics as composite bosons. These dimers can then form a Bose-Einstein condensate, giving this side of the crossover the designation of the BEC regime.

On the other side of the crossover ($a < 0$), no such two-particle bound state exists. However, the negative scattering length results in a weak attractive interatomic interaction. We will see that even for arbitrarily weak attraction, the Fermi gas creates a framework for which pairing can occur, a phenomenon called “Cooper pairing” [14]. Given that the full many-body description of a Cooper-paired Fermi gas has first been formulated by Bardeen, Cooper and Schrieffer, this side of the crossover is called the BCS regime [13].

Between these two regimes lies the unitarity regime, characterised by resonant scattering and very strong interactions. Here, the scattering length diverges ($a \rightarrow \infty$), causing it to drop out of all thermodynamic descriptions, leaving the interatomic spacing ($\propto 1/k_F$) as the only relevant length scale [8, 104].

As we will see, the transition from the regimes of composite dimers (BEC) to the many-body effect of Cooper pairing (BCS) is smooth, without a phase transition, signalling a crossover that smoothly transitions from fermionic to bosonic pairing [15, 104, 117–119]. In this crossover, we will explore the onset of superfluidity, a key aspect that will later be examined using machine learning techniques in Chapter 5.

2.3.1 BEC regime - Repulsive composite dimers

On the BEC side of the crossover, two fermions with distinguishable spin can form a composite dimer with a bound state energy as given in Eq. (2.10). If the bound state energy is much larger than all other energy scales of the gas, the dimers can be considered permanent composite bosonic molecules, composed of two spin 1/2 fermions. This statement is equivalent to the characteristic length scale of the molecules (a) being much smaller than the interatomic distance $1/(k_F a) \gg 1$. The parameter $1/(k_F a)$ will remain the most important parameter describing the thermodynamics throughout the crossover [8, 104].

In the BEC limit ($1/(k_F a) \rightarrow \infty$), the gas is made up of tightly-bound molecules of mass $m_B = 2m_{Li}$ and a molecular density equal to the original density per spin state $n_B = n_\sigma = n_{\sigma_1 + \sigma_2}/2$. Here, the critical temperature marking the onset of superfluidity equals the critical temperature of Bose-Einstein condensation of the composite dimers. The critical

temperature reads [8, 48, 103]

$$\begin{aligned}
 T_{\text{C,BEC}} &= \frac{\pi \hbar^2}{m_{\text{Li}}} \left(\frac{n_\sigma}{\zeta\left(\frac{3}{2}\right)} \right)^{\frac{2}{3}} \\
 \Leftrightarrow \frac{T_{\text{C,BEC}}}{T_{\text{F}}} &= \frac{2\pi}{\left(6\pi^2 \zeta\left(\frac{3}{2}\right)\right)^{\frac{2}{3}}} \approx 0.218,
 \end{aligned} \tag{2.12}$$

where, in the last step, the Fermi energy $E_{\text{F}} = k_{\text{B}}T_{\text{F}}$ from an ideal Fermi gas (Eq. (2.4)) cancels the density. Here, ζ denotes the Riemann zeta function.

While an ideal Bose gas undergoes Bose-Einstein condensation, the simultaneous occurrence of superfluidity requires interactions between the dimers to lift the critical velocity [103]. The scattering length a , as calculated from the Feshbach resonance, only directly describes the interatomic interaction. Due to the Pauli exclusion principle between the dimers' constituents, interactions between dimer-dimer and atom-dimer channels are drastically attenuated.

The effective atom-dimer scattering length has been calculated by [120] to be $a_{\text{ad}} = 1.18a$. For interactions between two composite dimers, the effective scattering length has been determined by [121], resulting in $a_{\text{dd}} = 0.6a$.

For stronger interactions, a first-order correction to the critical temperature $T_{\text{C,BEC},1}$ has been determined by [42, 43, 122–124]. This correction increases the critical temperature for stronger interactions approaching unitarity and is given by

$$\frac{T_{\text{C,BEC},1}}{T_{\text{C,BEC}}} = 1 + 1.31 \left(n_\sigma a_{\text{dd}}^3 \right)^{\frac{1}{3}} = 1 + \frac{1.31 \cdot 0.6}{\left(6\pi^2\right)^{\frac{1}{3}} \frac{1}{k_{\text{F}}a}}. \tag{2.13}$$

Generally, we find that bosonic physics begin to dominate around $1/(k_{\text{F}}a) \approx 0.6$ [125], where the single-particle chemical potential μ becomes negative.

2.3.2 BCS regime - Cooper pairing

The other end of the crossover is the BCS limit, which is reached for $1/(k_{\text{F}}a) \rightarrow -\infty$. In this regime, we consider a gas of weakly attractively interacting fermions in a two-state mixture. Unlike in the BEC regime, this regime does not support a two-body bound state; instead, pairing occurs due to the presence of the Fermi sea, even for arbitrarily small attractions V_0 .

The many-body Hamiltonian describing the balanced Fermi gas in the spin states $\sigma = \{\uparrow, \downarrow\}$ within the grand-canonical description ($\mu\hat{N}$) is given by [8, 103]

$$\hat{H} - \mu\hat{N} = \sum_{\mathbf{k}} \xi_{\mathbf{k}} \left(\hat{c}_{\mathbf{k},\uparrow}^\dagger \hat{c}_{\mathbf{k},\uparrow} + \hat{c}_{\mathbf{k},\downarrow}^\dagger \hat{c}_{\mathbf{k},\downarrow} \right) + V_0 \sum_{\mathbf{k}_1, \mathbf{k}_2, \mathbf{q}} \left(\hat{c}_{\mathbf{k}_1+\mathbf{q},\uparrow}^\dagger \hat{c}_{\mathbf{k}_2-\mathbf{q},\downarrow}^\dagger \hat{c}_{\mathbf{k}_2,\downarrow} \hat{c}_{\mathbf{k}_1,\uparrow} \right), \tag{2.14}$$

where $\xi_{\mathbf{k}} = \frac{\hbar^2 \mathbf{k}^2}{2m} - \mu$ represents the single-particle dispersion relation, absorbing the chemical potential μ . The first term of the Hamiltonian (Eq. (2.14)) is the sum over the single-particle kinetic energies across all momentum states \mathbf{k} . Interactions between fermions are considered in the second term.

As discussed in Section 2.2.1, only s -wave scattering is considered in these experiments, limiting scattering to distinguishable spin states σ . Additionally, since the interatomic distance cannot resolve the scattering potential, the generally momentum-dependent potential $V_{\mathbf{q}}$ can be replaced by a constant V_0 . Both of these simplifications are already incorporated in the second term, which describes scattering between two fermions with momenta $(\mathbf{k}_1, \mathbf{k}_2)$, transferring a momentum of \mathbf{q} , summed over all combinations of these three parameters, and scaled by the effective interatomic potential V_0 .

The general Hamiltonian in Eq. (2.14) can further be simplified by considering the involved momenta in the interaction term. In the Cooper pairing formalism, the bound state energy is proportional to the density of states. As two Fermions with net-zero momentum $\mathbf{k}_1 = -\mathbf{k}_2$ have the largest amount of final scattering states, scattering between two such constituents leads to the smallest possible energy. This leads to the further simplified Hamiltonian

$$\hat{H} - \mu\hat{N} = \sum_{\mathbf{k}} \xi_{\mathbf{k}} \left(\hat{c}_{\mathbf{k},\uparrow}^\dagger \hat{c}_{\mathbf{k},\uparrow} + \hat{c}_{\mathbf{k},\downarrow}^\dagger \hat{c}_{\mathbf{k},\downarrow} \right) + V_0 \sum_{\mathbf{k},\mathbf{q}} \left(\hat{c}_{\mathbf{k}+\mathbf{q},\uparrow}^\dagger \hat{c}_{-\mathbf{k}-\mathbf{q},\downarrow}^\dagger \hat{c}_{-\mathbf{k},\downarrow} \hat{c}_{\mathbf{k},\uparrow} \right). \quad (2.15)$$

As already mentioned, the bound state energy of Cooper pairs is proportional to the density of states $\rho(E)$ [8, 14, 103]. Unfortunately, in a 3D geometry ($\rho \propto \sqrt{E}$), an arbitrarily small attractive potential does not yield a bound state. However, the existence of the fully occupied Fermi sea prohibits momentum states below the Fermi energy E_F to participate in scattering, effectively restricting interactions to the region of the Fermi edge. Consequently, the density of states is replaced by its constant value at the Fermi energy $\rho(E) \rightarrow \rho(E_F)$. A constant density of states is also known from 2D geometries, where a bound state even exists for arbitrarily small potentials [8]. Thus, the existence of Cooper pairing in a weakly attractively interacting Fermi gas is a pure many-body effect.

In order to solve the Hamiltonian in Eq. (2.15) a mean-field approach is employed. Here, the pair-annihilation operator $\hat{b}_{\mathbf{k}} := \hat{c}_{-\mathbf{k},\downarrow} \hat{c}_{\mathbf{k},\uparrow}$ is introduced. In the mean-field approach, the pair-annihilation operator is split into its expectation value plus fluctuations $\hat{b}_{\mathbf{k}} = \langle \hat{b}_{\mathbf{k}} \rangle + \delta\hat{b}_{\mathbf{k}}$, where $\delta\hat{b}_{\mathbf{k}} = \hat{c}_{-\mathbf{k},\downarrow} \hat{c}_{\mathbf{k},\uparrow} - \langle \hat{b}_{\mathbf{k}} \rangle$. Finally, the decomposition is plugged back into the Hamiltonian in Eq. (2.15). By omitting all terms quadratic in fluctuation $\delta\hat{b}_{\mathbf{k}}$ and by introducing the new parameter $\Delta := -V_0 \sum_{\mathbf{k}} \langle \hat{b}_{\mathbf{k}} \rangle$ we arrive at the final BCS mean-field Hamiltonian, which reads

$$\hat{H} - \mu\hat{N} + |\Delta|^2/V_0 = \sum_{\mathbf{k}} \xi_{\mathbf{k}} \left(\hat{c}_{\mathbf{k},\uparrow}^\dagger \hat{c}_{\mathbf{k},\uparrow} + \hat{c}_{\mathbf{k},\downarrow}^\dagger \hat{c}_{\mathbf{k},\downarrow} \right) - \left(\Delta^* \hat{c}_{-\mathbf{k},\downarrow} \hat{c}_{\mathbf{k},\uparrow} + \Delta \hat{c}_{\mathbf{k},\uparrow}^\dagger \hat{c}_{-\mathbf{k},\downarrow}^\dagger \right). \quad (2.16)$$

The Hamiltonian now comprises diagonal terms describing the kinetic energy and off-diagonal terms stemming from the creation ($\hat{b}_{\mathbf{k}}^\dagger$) and annihilation ($\hat{b}_{\mathbf{k}}$) of Cooper pairs.

In order to solve the Hamiltonian, the Hamiltonian is diagonalised by a Bogoliubov transformation. This transformation introduces new annihilation and creation operators $\hat{\gamma}_{\mathbf{k},\uparrow}$ and $\hat{\gamma}_{-\mathbf{k},\downarrow}^\dagger$ as mixtures of the original fermionic operators $\hat{c}_{\mathbf{k},\uparrow}^\dagger$ and $\hat{c}_{-\mathbf{k},\downarrow}$ with amplitudes $u_{\mathbf{k}}$ and $v_{\mathbf{k}}$ in the form of

$$\begin{aligned} \hat{\gamma}_{\mathbf{k},\uparrow} &= u_{\mathbf{k}} \hat{c}_{\mathbf{k},\uparrow} - v_{\mathbf{k}} \hat{c}_{-\mathbf{k},\downarrow}^\dagger \\ \hat{\gamma}_{-\mathbf{k},\downarrow}^\dagger &= v_{\mathbf{k}} \hat{c}_{\mathbf{k},\uparrow} + u_{\mathbf{k}} \hat{c}_{-\mathbf{k},\downarrow}^\dagger. \end{aligned} \quad (2.17)$$

Interestingly, the new quasi-particles are linear combinations of particle and hole operators. Hence, quasi-particle excitations are coherent superpositions of particle and hole states. The amplitudes u_k, v_k of the mixing are given by

$$\begin{aligned} u_k^2 &= \frac{1}{2} \left(1 + \frac{\xi_k}{E_k} \right) \\ v_k^2 &= \frac{1}{2} \left(1 - \frac{\xi_k}{E_k} \right), \end{aligned} \quad (2.18)$$

with the dispersion relation of the quasi-particles being

$$E_k = \sqrt{\xi_k^2 + \Delta^2}. \quad (2.19)$$

Only here it becomes apparent that the previously introduced parameter Δ now functions as an energy gap in the new dispersion relation. No excitation can be created for energies $E < \Delta$. As quasi-particles must always be created in pairs to preserve the atom number, the minimum excitation energy is 2Δ . The gap has enormous consequences, as it gives the BCS state its superfluidity.

Finally, the ground state wave function of the BCS state $|\psi_{\text{BCS}}\rangle$ can be constructed from the vacuum $|0\rangle$ by minimising the total energy, which is given for zero occupation of quasi-particles. The ground state wave function therefore obeys $\hat{\gamma}_{\mathbf{k},\sigma} |\psi_{\text{BCS}}\rangle = 0$ and can be written as

$$|\psi_{\text{BCS}}\rangle = \prod_{\mathbf{k}} \left(u_k + v_k \hat{c}_{\mathbf{k},\uparrow}^\dagger \hat{c}_{-\mathbf{k},\downarrow}^\dagger \right) |0\rangle. \quad (2.20)$$

Here it becomes clear that $|u_k|^2$ denotes the probability to find a vacant pair state and $|v_k|^2$ gives the probability to find a pair at momentum \mathbf{k} , which thus obey $|u_k|^2 + |v_k|^2 = 1$.

By applying the original definition of the pairing gap $\Delta := -V_0 \sum_{\mathbf{k}} \langle \hat{b}_{\mathbf{k}} \rangle$ on the ground state $|\psi_{\text{BCS}}\rangle$, we find the following relation

$$\Delta = -V_0 \sum_{\mathbf{k}} \langle \psi_{\text{BCS}} | \hat{b}_{\mathbf{k}} | \psi_{\text{BCS}} \rangle = -V_0 \sum_{\mathbf{k}} u_k v_k. \quad (2.21)$$

First, we find that the gap parameter Δ is proportional to $\sum_{\mathbf{k}} u_k v_k$. Hence, only pairs with $u_k, v_k \neq 0$ contribute to the gap Δ and can be considered part of the condensate. Thus, the joint expression $u_k v_k$ can also be interpreted as the condensed macroscopic wave function of Cooper pairs [8, 69]. The number of condensed pairs is actually defined as $N_0 = \sum_{\mathbf{k}} |u_k v_k|^2$ [126].

In order to get an understanding of pairing in terms of u_k, v_k and E_k throughout the crossover, we have to determine both the gap Δ and chemical potential μ as function of the interaction parameter $1/(k_F a)$. By solving both the gap equation Eq. (2.21) and the number equation $N \propto \int d\mathbf{k}^3 v_k^2$ self consistently, results of the mean-field BCS approach can be calculated. This result is known as the mean-field result, and even though no easy analytical expression can be obtained, a numerical solution is shown in Fig. 2.5. However, it has been shown that an analytical result can be obtained by elliptic integrals [127]. The figure also shows values of the chemical potential μ and gap Δ in the limits of the BEC

regime $1/(k_F a) \rightarrow \infty$ and BCS regime $1/(k_F a) \rightarrow -\infty$. In these limits, the mean-field result provides approximate analytical results.

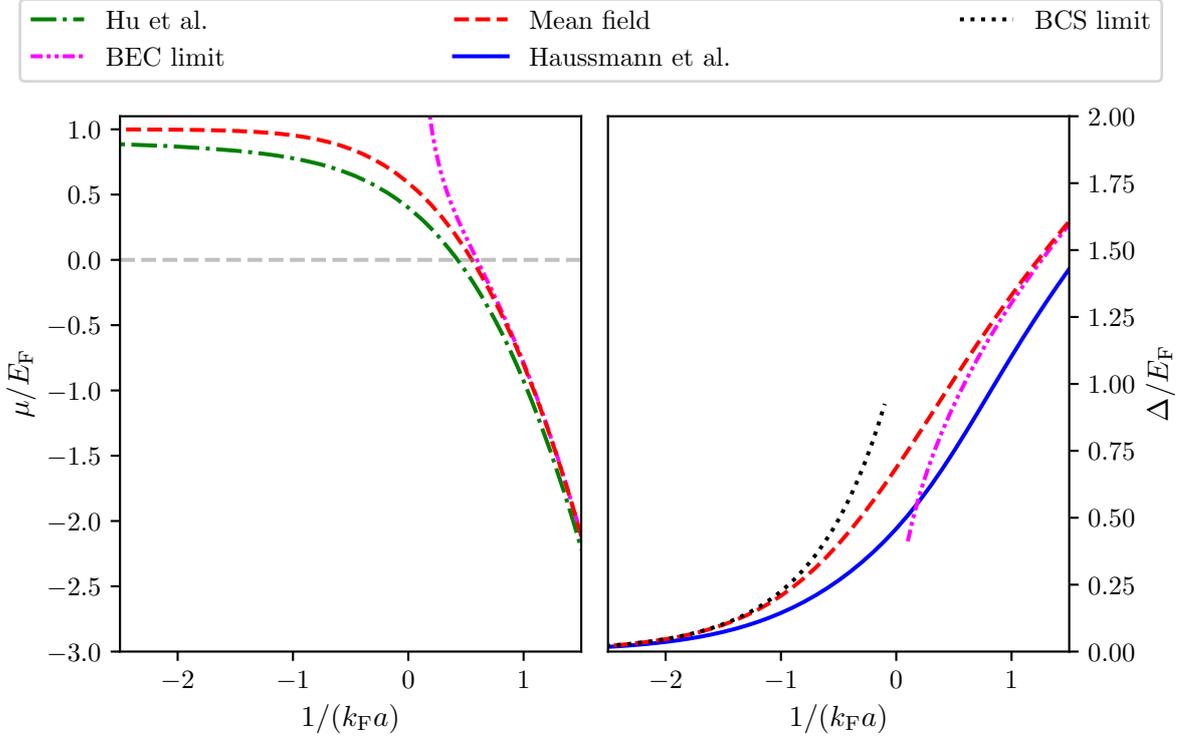


Figure 2.5: Chemical potential μ and gap Δ in the BEC-BCS crossover. The figure presents several theoretical models for a homogeneous gas at zero temperature $T = 0$. The chemical potential μ considers an emerging bound state $E_B/2 = (k_F a)^{-2}$ on the BEC side of the crossover. The more intricate results from Hu et al. and Haussmann et al. are reproduced from [125] and [45], respectively. The mean-field result is calculated via the mean-field approach in Section 2.3.2 and can also be found in [8]. Additionally, the analytical results in the BEC and BCS limit are shown, which can be found in Eq. (2.23) and Eq. (2.22), respectively.

In the BCS Limit, the chemical potential μ and gap Δ read (**BCS Limit**) [8, 103]

$$\begin{aligned} \mu &\approx E_F \\ \Delta &\approx \frac{8}{e^2} e^{\pi/(2k_F a)}. \end{aligned} \quad (2.22)$$

As the available interaction energy cannot soften the Fermi edge very much and Pauli blocking still dominates, the chemical potential is close to the Fermi energy. This scenario resembles an ideal Fermi gas very closely (compare to Fig. 2.1), which is a special case of the BCS wavefunction $|\psi_{\text{BCS}}\rangle$ for $u_k = 0$. In such a system, the gap is exponentially suppressed being proportional to the expectation value of the pair annihilation operator $\langle \hat{b}_{\mathbf{k}} \rangle = u_k v_k$ (Eq. (2.21)). Basically, the Fermi sea only melts in a narrow region of Δ around the Fermi edge, leading to fragile pairing.

In the BEC Limit, the chemical potential μ and gap Δ read (**BEC Limit**) [8, 103]

$$\begin{aligned}\mu &\approx -\frac{\hbar^2}{2m_{\text{Li}}a^2} + \frac{2\pi\hbar^2 an_\sigma}{m_{\text{Li}}} \\ \Delta &\approx \sqrt{\frac{16}{3\pi}} \frac{E_{\text{F}}}{\sqrt{k_{\text{F}}a}}.\end{aligned}\tag{2.23}$$

Here, the chemical potential comprises two terms. The first term is the bound state energy (per atom) of the tightly-bound molecules, as introduced in Eq. (2.10). As previously discussed in Section 2.3.1, the composite dimers also interact with each other, leading to a repulsive ($a > 0$) mean-field contribution in the second term. While this term closely resembles the chemical potential of a condensed Bose gas [48], it lacks the suppressed dimer-dimer scattering length of $a_{\text{dd}} = 0.6$ [121]. Later in this thesis, the repulsive mean-field contribution in a composite Bose gas becomes important in enabling a novel cooling method, explored in Chapter 7.

As discussed beforehand, the chemical potential μ is mostly governed by the negative bound state energy in the BEC limit. Hence, the gap parameter Δ cannot describe the molecular bound state, nor does it correspond to the minimal excitation energy $\min(E_k)$ in Eq. (2.19). Instead, the minimum in the dispersion relation is found at $k = 0$ for $\mu < 0$ and can be approximated to

$$E_k = \sqrt{\xi_k^2 + \Delta^2} \approx |\mu| + \frac{\Delta^2}{2|\mu|}.\tag{2.24}$$

Here, the minimum excitation energy is lifted from the single-particle bound state $|\mu|$ by the pairing mean-field $\Delta^2/(2|\mu|)$, experienced by unpaired fermions in the molecular gas.

Both the results in the BEC limit (Eq. (2.23)) and BCS limit (Eq. (2.24)), as well as the full mean-field result are shown in Fig. 2.5. Additionally, more intricate theoretical results for the chemical potential μ [125], considering pair fluctuations, and the gap parameter Δ [45], based on a variational many-body approach, are presented here as well.

Following the smooth evolution of both the chemical potential μ and gap Δ within the strongly interacting regime $-1 < 1/(k_{\text{F}}a) < 1$ (Fig. 2.5) already suggests why there is no phase transition, but rather a crossover. In the limit of weakly attractive fermions, the chemical potential closely resembles that of an ideal Fermi gas at low temperatures (compare to Fig. 2.1). As interactions become stronger, the Fermi sea gradually loses its step-function-like behaviour, allowing lower-energy states to be populated, which in turn reduces the chemical potential. Further into the resonance, the chemical potential flips sign around $1/(k_{\text{F}}a) \approx 0.6$. Here, the physics become increasingly bosonic, progressing towards the BEC limit $1/(k_{\text{F}}a) \rightarrow \infty$, where we have tightly-bound molecules.

This same behaviour is also reflected in the size of the pairs [17]. In the BCS limit, pairs are significantly larger than the interparticle spacing. Near unitarity, the pair size gradually shrinks down to the scale of the interparticle spacing. Finally, in the BEC limit, we observe tightly-bound molecules.

Using the results for the chemical potential μ and gap Δ , we can now understand the pairing behaviour throughout the crossover. The parameters $|u_k|^2$ and $|v_k|^2$ in Eq. (2.18) represent the non-occupation and occupation of a momentum state \mathbf{k} with a pair, respectively.

In Fig. 2.6, we show these parameters along with the mixed quantity $|u_k v_k|^2$, which represents the macroscopic and condensed Cooper pairing wave function. We show all three quantities for three exemplary interaction parameters $1/(k_F a) = \{-1.0, 0.0, 1.0\}$, using the results of the mean-field approach (refer to Fig. 2.5). Additionally, we also show the corresponding quasi-particle dispersion relation from Eq. (2.19).

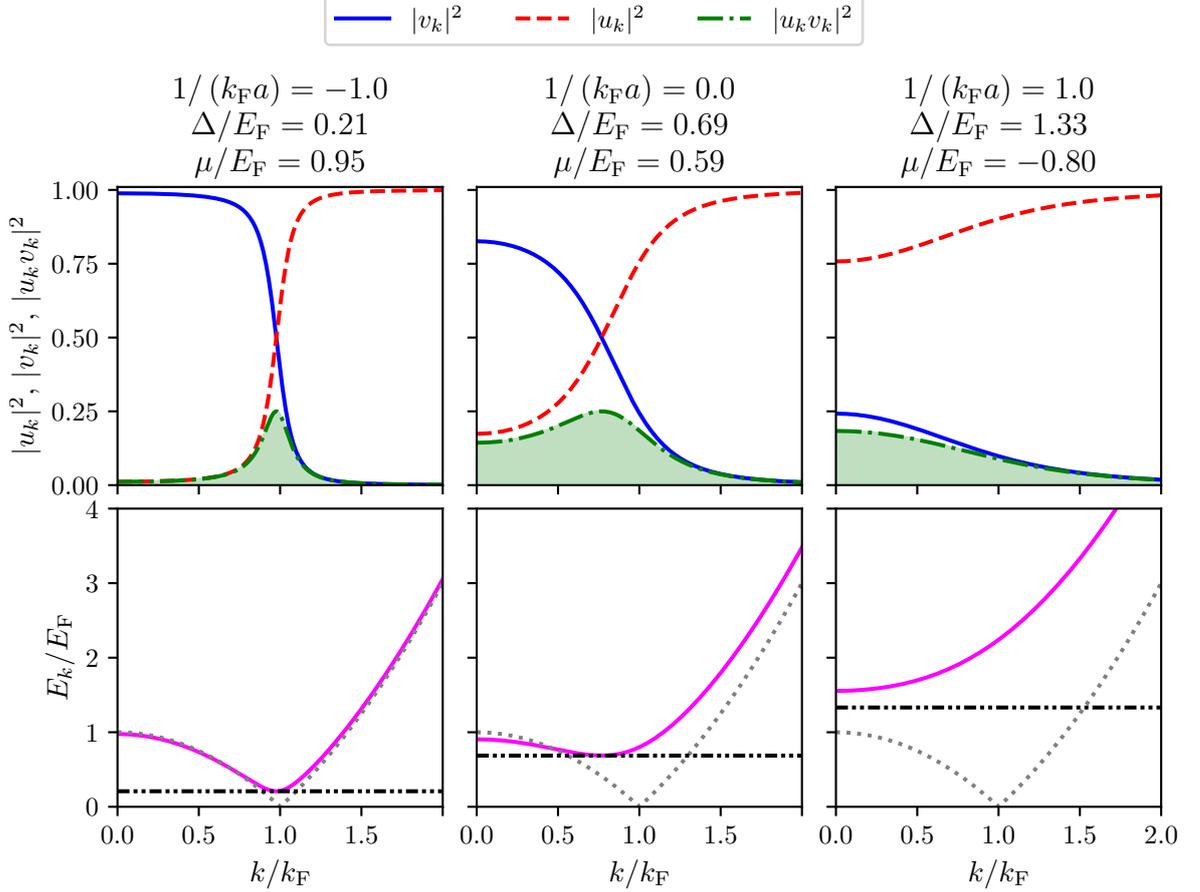


Figure 2.6: Momentum distribution and dispersion relation in the BEC-BCS crossover (mean-field results) at zero temperature. **Upper row:** The momentum distribution of the parameters $|u_k|^2$ and $|v_k|^2$ (Eq. (2.18)) is shown throughout the crossover, corresponding to the non-occupation and occupation with pairs at the momentum state \mathbf{k} , respectively. Above the plots, the respective interaction parameters $1/(k_F a)$ and mean-field results of the gap Δ and chemical potential μ are given, compare to Fig. 2.5. Additionally, the mixed quantity $|u_k v_k|^2$ is shown. The shaded area is proportional to the gap Δ , according to Eq. (2.21). **Lower row:** The corresponding quasi-particle dispersion relation (solid line) (refer to Eq. (2.19)) at the respective interaction parameters is illustrated, along with the dispersion relation of the ideal Fermi gas (dotted line) and the mean-field value of the gap Δ (dash-dot-dot).

Well on the BCS side ($1/(k_F a) = -1$), the Fermi sea $|v_k|^2$ is almost intact, with just a few momentum states below the Fermi energy E_F being empty. Here, Cooper pairing concentrates along the Fermi edge, which is visible in the peaked structure of the mixed parameter $|u_k v_k|^2$. While the gap Δ is small, it is non zero, lifting the minimum in the dispersion relation to

$\min(E_k) = \Delta$. The minimum in the dispersion relation is close to the Fermi momentum at $\xi_k + \mu \approx E_F$.

Moving to unitarity ($1/(k_F a) = 0$), the Fermi sea melts down more drastically. Here, no momentum state \mathbf{k} has close to full occupation, and so the macroscopic wave $|u_k v_k|^2$ function extends over all momentum states with a small peak towards larger momenta $k < k_F$. As the pairing volume increases, so does the gap Δ . While the minimum position in the dispersion relation shifts to smaller momenta $k < k_F$, the chemical potential is still positive $\mu > 0$, and so a minimum can be found at $\min(E_k) = \Delta$.

Approaching the BEC regime ($1/(k_F a) = 1$), the chemical potential turns negative $\mu < 1$ as molecular binding starts to dominate. Here, pairing occurs over all momentum states with a peaked structure at zero momentum $k = 0$. Therefore, Δ no longer depicts the minimum excitation energy but an additional pairing field, which lifts the excitation spectrum at its minimum position $k = 0$, as discussed in Eq. (2.24).

2.3.3 Superfluid phase transition

As we have seen in Section 2.3.1 and Section 2.3.2, pairing persists throughout the entire crossover. This has been demonstrated by the smooth evolution of the pairing coefficients u_k and v_k in Fig. 2.6 within the strong coupling regime $-1 < 1/(k_F a) < 1$. However, as the pairing mechanisms transition from fermionic to bosonic in nature, the critical temperature T_C of the superfluid phase transition no longer aligns with the onset of pairing. Calculating the critical temperature across the crossover is very much complicated due to the strong interactions around unitarity. By examining the crossover limits, however, we can gain an initial sense of the critical temperature trend.

In the BEC limit, weakly repulsive composite bosonic dimers can undergo Bose-Einstein condensation. Here, the critical temperature of the superfluid phase transition coincides with the critical temperature of condensation as stated in Eq. (2.12). Additionally, the first-order correction to the critical temperature in Eq. (2.13) shifts the critical temperature upwards with stronger interactions.

In the BCS limit, we have discussed that even for arbitrarily-small attractive coupling, a finite gap Δ (Eq. (2.22)) opens in the excitation spectrum in Eq. (2.19). Similarly, the critical temperature in this limit can be obtained by solving the mean-field approach for a vanishing gap Δ . The critical temperature in the BCS limit reads [44]

$$\frac{T_{C,\text{BCS+GMB}}}{T_F} = \frac{1}{(4e)^{\frac{1}{3}}} \frac{e^\gamma}{\pi} \frac{8}{e^2} e^{\pi/(2k_F a)}, \quad (2.25)$$

where γ is the Euler constant, obeying $e^\gamma \approx 1.78$ [8]. In comparison to the gap in the BCS limit in Eq. (2.22), the critical temperature is scaled by an additional factor of e^γ/π . Additionally, the stated critical temperature also includes a beyond mean-field correction, which reduces the critical temperature by another factor of $1/(4e)^{1/3} \approx 0.45$ [44]. The correction is known as the Gorkov-Melik-Barkhudarov (GMB) correction, and considers screening of the interparticle interaction due to pairing polarisation effects within the fermionic sample [44].

We see that both the critical temperature in the BEC limit with its first-order correction (Eq. (2.13)), as well as the critical temperature in the BCS limit (Eq. (2.25)) increase for

stronger interactions. This simple argument demands that a maximum in critical temperature must exist somewhere in the unitarity regime. Both results of the critical temperature in the BEC and BCS limit are shown in Fig. 2.7. Additionally, several theoretical predictions of the phase boundary are shown. Measurements of the critical temperature have been performed during this thesis and are presented later in Chapter 5.

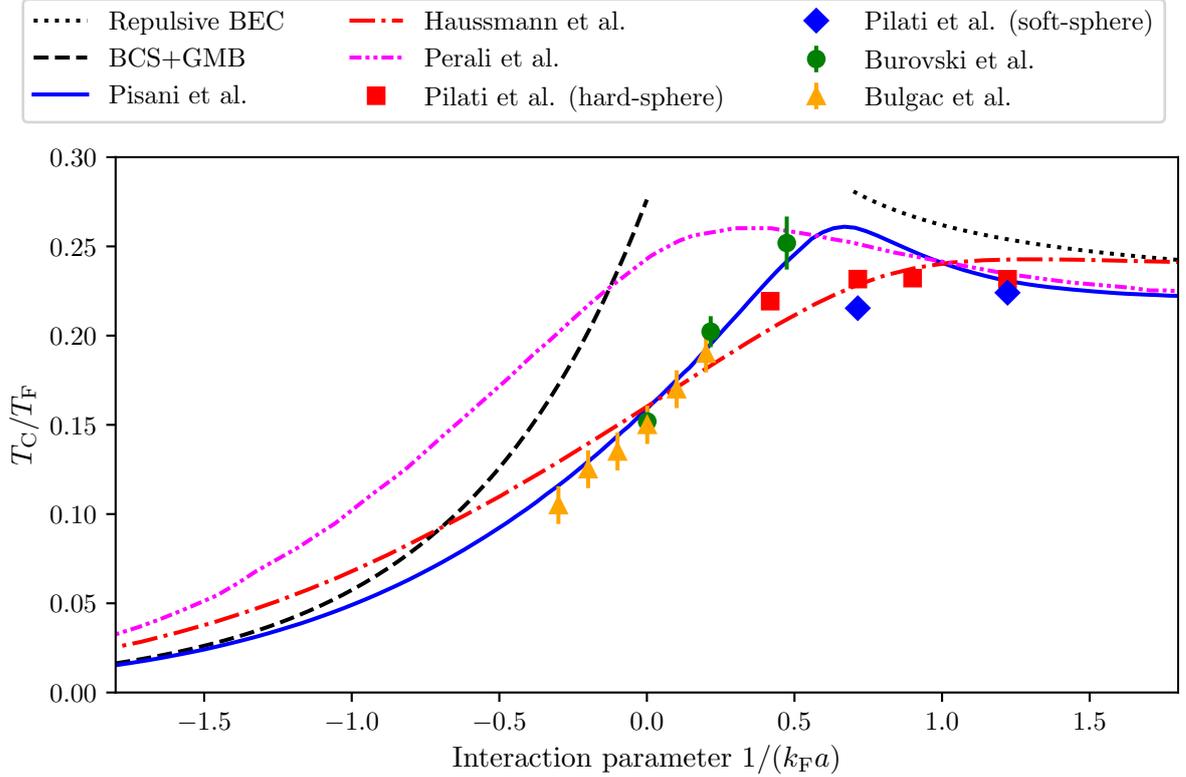


Figure 2.7: Critical Temperature in the BEC-BCS crossover. The figure presents several theoretical calculations of the critical temperature T_C/T_F . In the limits, the analytical results for the weakly interacting BEC (Eq. (2.13)) (dotted line) and the exponentially decaying BCS result with GMB correction (Eq. (2.25)) (dashed line) are shown. Across unitarity, an extended GMB correction [19] (solid line) and a self-consistent t-matrix approach [45] (dashed-dotted line) are plotted. Additionally, two quantum Monte Carlo results (circles [128] and triangles [47]) are presented. Finally, quantum Monte Carlo results for hard-sphere (red square) and soft-sphere (blue diamond) bosons are shown [20].

2.3.4 Collective excitations

In Section 2.3, we have introduced the fundamental quasi-particle dispersion relation in Eq. (2.19). This excitation spectrum describes a “pair-breaking” mechanism through the creation of an unpaired particle and a hole, thus requiring two quasi-particle excitations to conserve atom number. In the BCS limit, the excitation minimum occurs near the Fermi edge $k \approx k_F$ and requires an energy of 2Δ . Towards the molecular BEC limit, the lowest excitation occurs at zero momentum $k = 0$ and is raised by a mean-field shift experienced by

the unpaired fermions in the molecular sample, as described in Eq. (2.24). Since pair-breaking concerns single particles, this spectrum is known as the “single-particle excitation spectrum”.

However, in the considered many-body state, another class of excitations can occur, namely “collective excitations”. A particularly interesting class of collective excitations arises from the second-order phase transition associated with the superfluid phase transition, as described by the Ginzburg-Landau theory [70]. Here, the previously introduced gap parameter Δ of the superfluid acts as a complex order parameter. As previously discussed, the order parameter not only describes the superfluid gap Δ (Eq. (2.19)) but also the macroscopic wave function of the Cooper pairs responsible for superfluidity (Eq. (2.21)) [126].

The Ginzburg-Landau theory shows that for such an order parameter Δ , the free energy F of the system can be parametrised as [70]

$$F = F_0 + \alpha|\Delta|^2 + \frac{\beta}{2}|\Delta|^4, \quad (2.26)$$

where the composition of the real parameters α and β change according to the phase of the system. Above the phase transition $T > T_C$, both parameters are positive $\Rightarrow \alpha, \beta > 0$. Thus, the system is in its normal phase and the ground state can be found in the centre of the complex plane at $\Delta = 0$.

Once the phase transition is crossed $T \leq T_C$, the new ground state can be found at a finite value of the order parameter $|\Delta_0| > 0$. Here, the parameters of the Ginzburg-Landau free energy take values obeying $\alpha < 0$ and $\beta > 0$. Thus, the ground state is described by $|\Delta_0|^2 = -\alpha/\beta$. As the order parameter Δ is a complex quantity, it can be described by an amplitude and a phase $\Delta_0 = \sqrt{|\Delta_0|^2} e^{i \arg(\Delta_0)} = \Re\Delta_0 + i\Im\Delta_0$. The phase can be interpreted as the phase of the macroscopic wave function. Since the phase does not alter the system’s Hamiltonian, it is arbitrarily chosen during the phase transition, effectively breaking the system’s $U(1)$ symmetry [70]. This scenario is sketched in Fig. 2.8, with the ground state highlighted by the red ball at $\Delta = \Delta_0$. Due to the shape of the free energy in the complex plane, the free energy is often referred to as the “mexican hat” potential.

In a naive picture, the shape of the potential suggests two distinct excitation modes. First, the ball can be moved along the valley of the potential, effectively changing the phase of the order parameter. This excitation, depicted by the yellow-dashed arrow, is known as the (Nambu-)Goldstone mode or phase mode [71, 72]. As the energy does not change along this mode, arbitrarily small excitations can be performed, which is why this mode is also called massless. These excitations manifest as sound modes, known as Bogoliubov-Anderson modes in a BCS system [129].

More complicated, the ball can also be moved radially, changing the amplitude of the order parameter Δ_0 . This excitation is highlighted by the green solid arrow and is known as the Higgs mode or amplitude mode [70]. Similarly to the harmonic oscillator, the order parameter experiences a harmonic pseudo potential along this direction, resulting in a gapped excitation that requires a minimum energy of $2\Delta_0$. Exciting this mode induces oscillations of the order parameter with a frequency of $\omega_H = 2\Delta/\hbar$ [73].

With a higher energy than the gapless phase mode, the Higgs mode is generally unstable and decays into lower-energy sound excitations. On the BCS side, the mode is considered more stable and has a characteristic power law damping of $t^{-1/2}$ [73]. Stability in this regime

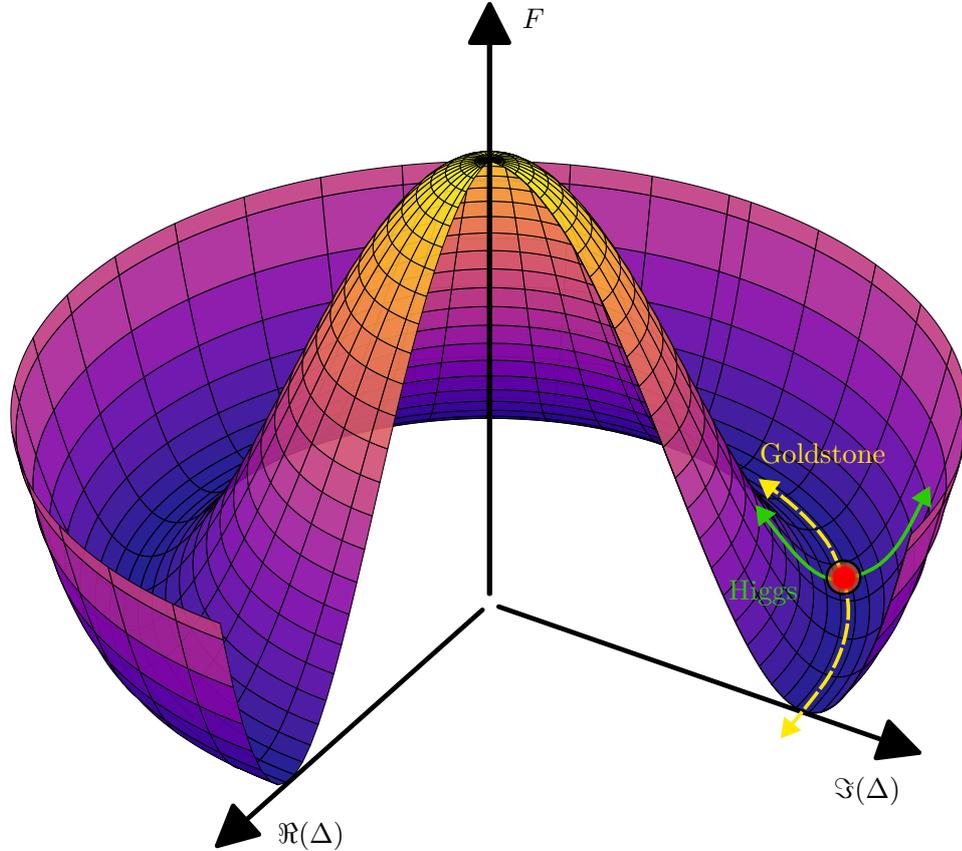


Figure 2.8: Mexican hat potential below a superfluid phase transition. Crossing the phase transition, the free energy F (Eq. (2.26)) landscape forms a ring-like potential well in the complex plane at $|\Delta| = \Delta_0$, defining the ground state energy. Since $\Delta = |\Delta| \exp^{i\varphi}$ is a complex quantity, a particular phase φ is chosen (red dot), breaking the system's symmetry. The energy landscape supports two fundamentally distinct excitation modes: the gapless Goldstone phase mode (dashed, yellow arrow) and the gapped Higgs amplitude mode (green, solid arrow). Figure inspired from [70].

arises from particle-hole symmetry, which introduces an effective Lorentz invariance in the system, allowing for a stable Higgs mode [71, 84].

Particle-hole symmetry is evident in the quasi-particle operators from Eq. (2.18) and their dispersion relation in Eq. (2.19). In the BCS state, the energy to add a single fermion to state (\mathbf{k}', \uparrow) is the sum of the single particle kinetic energy ξ_k plus the negative pairing energy $E_k - \xi_k$, which arises because the state and the mirrored partner state $(-\mathbf{k}, \downarrow)$ are now closed for the other particles. The total energy to add a single Fermion adds up to E_k and is reflected by the quasi-particle excitation $\hat{\gamma}_{k,\uparrow}$ [8]. Interestingly, the energy to remove a particle from state (\mathbf{k}', \uparrow) or to create a hole costs the same energy of E_k . In the BCS limit, the gap parameter is defined by the region around the Fermi edge. Here, the energy to create a particle- ($k > k_F$) or hole-excitation ($k < k_F$) close to the Fermi energy $k = k_F$ is

symmetric around the Fermi energy $\xi_k = 0$, because the gap Δ is exponentially suppressed [130]. This is known as particle-hole symmetry.

For stronger interactions, the gap Δ increases, and the Fermi sea softens over a broader region of $\sim \Delta$. Thus, particle- and hole-excitations gradually lose their symmetry. In the BEC limit, the damping is predicted to have a stronger power law dependence of $t^{-3/2}$ [131]. Additionally, the Higgs frequency shifts to $\omega_H = 2\sqrt{\Delta^2 + \mu^2}/\hbar$, as the smallest excitation for $\mu < 0$ can be found at zero momentum $k = 0$, as discussed in Eq. (2.24).

In this thesis, two distinct excitation methods of the Higgs mode are analysed in Chapter 6. Exciting the Higgs mode is non-trivial since the order parameter lacks typical coupling properties, such as an electric charge or a magnetic moment [132]. To induce perturbations effectively, the entire pseudo potential must be modulated [70]. This is achieved by perturbing the interaction parameter through a superimposed magnetic field. First, a parametric modulation of the interaction parameter is performed, allowing for a spectroscopic measurement of the Higgs mode. The modulation of the interaction parameter leads to an effective modulation of the order parameter, since both are coupled (refer to Fig. 2.5). Next, the order parameter is excited by a quench of the interaction parameter, which couples to all modes with a frequency lower than the ramp speed of the quench. Naturally, this requires a very fast change of the magnetic field (refer to Section 4.5.1). Afterwards, a time-resolved oscillation of the order parameter can be observed.

2.4 Fermions in an inhomogeneous trap

Up to this point, we have only considered fermions in a homogeneous system with zero potential energy. As we will see in Chapter 3, the fermions in this work are confined in an approximately harmonic 3D trapping potential. Such a trapping potential $U(\mathbf{r} = (x, y, z))$ is generally defined by three distinct trapping frequencies $\omega_i/(2\pi)$ along each Cartesian direction $i = \{x, y, z\}$ and can be expressed as

$$U(\mathbf{r}) = \frac{m}{2} \sum_{i=x,y,z} (\omega_i r_i)^2, \quad (2.27)$$

where m is the mass of the trapped particles, which, in this thesis, is the mass $m = m_{\text{Li}}$ of ${}^6\text{Li}$ atoms.

The potential can confine a certain number of fermions depending on the temperature T of the gas. In order to calculate the Fermi energy in terms of the atom number N_σ per spin state, we apply the ‘‘local density approximation’’ (LDA) [8, 103], similarly to the integral in Eq. (2.4).

In the LDA, the inhomogeneous potential energy $U(\mathbf{r})$ is absorbed into the chemical potential

$$\mu \rightarrow \mu(\mathbf{r}) = \mu - U(\mathbf{r}), \quad (2.28)$$

effectively creating an inhomogeneous chemical potential. Next, the Fermi-Dirac distribution

is integrated over the full phase-space at $T = 0 \Rightarrow \mu = E_F$, yielding the atom number

$$\begin{aligned}
 N_\sigma &= \int d^3\mathbf{r} \int \frac{d^3\mathbf{p}}{(2\pi\hbar)^3} f_{\text{Fermi}}\left(\frac{p^2}{2m} + U(\mathbf{r}), T = 0\right) \\
 &= 4\pi \int d^3\mathbf{r} \int_0^{\sqrt{2m(E_F - U(\mathbf{r}))}} \frac{dp}{(2\pi\hbar)^3} p^2 \\
 &= \int d^3\mathbf{r} \frac{1}{6\pi^2} \left(\frac{2m(E_F - U(\mathbf{r}))}{\hbar^2}\right)^{\frac{3}{2}} \\
 &= \frac{1}{6} \left(\frac{E_F}{\hbar\bar{\omega}}\right)^3,
 \end{aligned} \tag{2.29}$$

where $\bar{\omega} = \sqrt[3]{\prod_{i=x,y,z} \omega_i}$ is the geometric mean of the trapping frequencies. This results in a Fermi energy of [108]

$$E_F = \hbar\bar{\omega}(6N_\sigma)^{\frac{1}{3}}, \tag{2.30}$$

describing the sample in the harmonic trap.

In Eq. (2.4), we have calculated the Fermi energy $E_F = \frac{\hbar^2}{2m} (6\pi^2 n_\sigma)^{\frac{2}{3}}$ in terms of the local (homogenous) density n_σ . It is important to note that the Fermi energy of the harmonic trap only equals the local (homogeneous) Fermi energy calculated from the centre density n_σ (Eq. (2.4)) for an ideal Fermi gas without interactions. Introducing interactions can significantly alter the density profile, introducing deviations between both calculation of the Fermi energy. For instance, far on the BEC side, a Bose-enhanced density due to Bose-Einstein condensation may occur, changing the corresponding Fermi energy significantly. Therefore, we always use the Fermi energy $E_F = \frac{\hbar^2}{2m} (6\pi^2 n_\sigma)^{\frac{2}{3}}$ determined from the density n_σ at the centre of the trap (refer to Eq. (2.4)). This is particularly important when describing phenomena such as the onset of superfluidity, which typically first occurs at the location of maximum Fermi energy, and thus largest density.

The LDA is valid when the Fermi energy is much larger than the level spacing of the trap $E_F \gg \hbar\bar{\omega}$ [8, 103]. In our samples, this condition generally holds, with most samples exhibiting Fermi energies of approximately $E_F/h \approx 20$ kHz, while the level spacing is typically around $\omega/(2\pi) \approx 200$ Hz.

For a non-zero temperature $T > 0$ or different trapping potentials $U(\mathbf{r})$, the same LDA approach can still be used to calculate the normal state density in spatial coordinates $n_\sigma(\mathbf{r})$ or in momentum space $n_\sigma(\mathbf{p})$ by performing the integrals

$$\begin{aligned}
 n_\sigma(\mathbf{r}) &= \int \frac{d^3\mathbf{p}}{(2\pi\hbar)^3} f_{\text{Fermi}}\left(\frac{p^2}{2m} + U(\mathbf{r}), T\right) \\
 n_\sigma(\mathbf{p}) &= \int d^3\mathbf{r} f_{\text{Fermi}}\left(\frac{p^2}{2m} + U(\mathbf{r}), T\right).
 \end{aligned} \tag{2.31}$$

Apparatus and Preparation

This chapter provides an overview of the experimental apparatus used to prepare ultracold fermionic samples necessary for simulating the BEC-BCS crossover, as introduced in Section 2.3. The apparatus has initially been planned, designed and constructed by the first two PhD students in this experiment, Dr. Alexandra Behrle and Dr. Timothy Harrison, as well as by Dr. Kuiyi Gao. More detailed descriptions of the system can be found in the dissertations of the first two students [40, 41]. Details about minor upgrades and additional details are available in the later dissertations from this experiment [133–135]. In this thesis, the most important aspects to prepare ultracold samples of ${}^6\text{Li}$ are revisited and discussed.

While ${}^6\text{Li}$ presents our fermionic isotope used to simulate the BEC-BCS crossover, the apparatus also hosts a bosonic sample of sodium ${}^{23}\text{Na}$. Here, ${}^{23}\text{Na}$ solely serves as a coolant for sympathetic cooling of ${}^6\text{Li}$. Accordingly, this chapter begins by introducing both species' transition levels, which are exploited to manipulate the atoms with controlled laser light. As the atoms are cooled towards their ground state, the magnetic hyperfine splitting becomes important for magnetic trapping and for preparing a two-level mixture to simulate the crossover. Thus, we calculate and discuss the hyperfine magnetic sublevel structure of each species' ground state.

To create dilute and ultracold atomic samples, a chamber with ultra-high vacuum is necessary. Therefore, the next section presents the vacuum system of the experimental apparatus. We then discuss the path of the initially thermal mixture of isotopes towards an ultracold, degenerate fermionic sample of ${}^6\text{Li}$. In the end, the sample is held by an optical dipole trap. Here, we explore the degrees of control that constitute our quantum simulator.

Outline of this chapter

The sections of this chapter are organised as follows:

- Section 3.1: This section describes and discusses the transition levels and magnetic hyperfine splitting of the two species, ${}^6\text{Li}$ and ${}^{23}\text{Na}$, used in this experiment.
- Section 3.2: Here, we provide an overview of the performed steps to prepare a degenerate Fermi gas from an initially hot sample.

3.1 Properties of Lithium and Sodium

As previously mentioned, the apparatus contains a mixture of the fermionic isotope ${}^6\text{Li}$ and the bosonic isotope ${}^{23}\text{Na}$. In order to manipulate, control and detect these isotopes, the level structure is driven by resonant or off-resonant laser fields. This section will first introduce the relevant level structure for controlling the atoms with laser light. Next, we will discuss the magnetic sublevels of the hyperfine ground state. Properties about ${}^6\text{Li}$ and ${}^{23}\text{Na}$ are detailed in [136] and [137], respectively.

3.1.1 Transition levels and laser systems

Since both isotopes are alkalis, they possess a hydrogen-like level structure with a single valence electron of spin $S = 1/2$. Moreover, the ground state of either isotope is given by the hyperfine splitting of the state with zero angular momentum $L = 0 \Rightarrow J = S$. In hyperfine coupling, the electron's magnetic field ($\propto S$) couples to the magnetic moment of the nucleus I , manifesting in the total angular momentum $F = I \pm S$. Thus, the fine structure ground state splits into two hyperfine levels with $F = I \pm 1/2$.

Dipole transitions ($|\Delta L| = 1$) from the ground state couple to the $L = 1$ level, which has two fine structure sublevels of $J = 3/2$ and $J = 1/2$, each with its own hyperfine manifold. Transitions to the $J = 1/2$ manifold are known as the D_1 line, while transitions to the $J = 3/2$ manifold are known as the D_2 line. In this setup, we solely use transitions from the D_2 -line for both isotopes. The D_2 lines for both isotopes are illustrated in Fig. 3.1 alongside the splitting of the corresponding hyperfine manifolds.

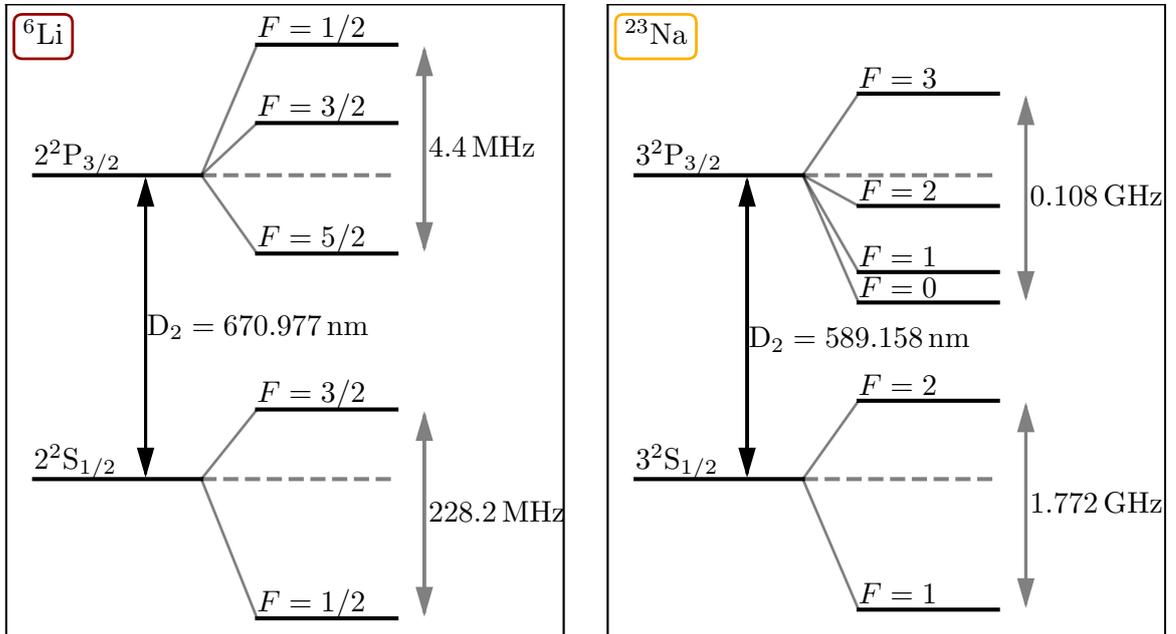


Figure 3.1: Level diagram of the D_2 -line transitions in ${}^6\text{Li}$ and ${}^{23}\text{Na}$. The figure shows the hyperfine structure in lithium (left) and sodium (right) of the corresponding ground state and ${}^2P_{3/2}$ excited state. Transition frequencies are taken from [136] for ${}^6\text{Li}$ and [137] for ${}^{23}\text{Na}$.

Lithium laser system

The carrier frequency of the lithium laser system is generated by a custom-built “external-cavity diode laser” (ECDL) operating at approximately $\lambda = 671$ nm. This laser is locked to an external reference vapour cell [138] via Doppler-free spectroscopy. The light is amplified by two tapered amplifiers (TA) in series, providing roughly 200 mW of power for the six arms of the magneto-optical trap (MOT), which is discussed further in Section 3.2.4.

A small fraction of light from the first laser is directed to a second ECDL, which is offset-locked to the first laser. The second laser, amplified by another TA, supplies light for the Zeeman slower, optical pumping and repumping. Additionally, light from this laser is used for imaging of the atomic sample. As we will explore in Section 3.1.2, the resonance frequency changes significantly with the external magnetic field. To keep the imaging on resonance across a scan range up to 1000 G of magnetic field, the offset lock is ramped prior to imaging by a magnitude of 1 GHz using an electro-optical modulator (EOM).

Several acousto-optical modulators (AOM) are used to adjust frequencies and power balancing e.g., to create offsets before locking or to optimise the experiments’ performance.

Sodium laser system

The light for sodium originates from a commercial* ECDL running at 1178 nm. Before frequency doubling, the light is amplified by a commercial† Raman fiber amplifier to 2.7 W. This amplified light seeds a custom-built frequency-doubling cavity [40], producing approximately 1.4 W of $\lambda = 589$ nm light. The cavity is locked via the Pound-Drever-Hall technique [139].

The cavity has two outputs. The primary output has higher power of ~ 1.2 W and is used for the six arms of the sodium MOT (total ~ 200 mW), as well as for the dark spot, imaging, pumping and repumping, which is all further explored in Section 3.2. The secondary output of the cavity has ~ 200 mW of power and is mainly used for the Zeeman slower, with a small fraction used for locking. Here, the laser frequency is locked via Doppler-free spectroscopy in another custom-built vapour cell containing sodium [140].

As with the lithium laser system, several AOMs are employed to adjust power and frequency of different beam paths, optimising the experiments’ performance.

3.1.2 Hyperfine splitting of the ground state

So far, we have only considered the degenerate hyperfine ground states of both ${}^6\text{Li}$ ($2^2\text{S}_{1/2}$, $F = 1/2$ and $F = 3/2$) and ${}^{23}\text{Na}$ ($3^2\text{S}_{1/2}$, $F = 1$ and $F = 2$). When an external magnetic field is applied, this degeneracy is lifted, and a manifold of $2F + 1$ distinct magnetic sublevels emerges. As discussed in Section 2.2.2, we use strong magnetic fields to tune interactions in a state mixture. These states are represented by the magnetic sublevels.

Magnetic sublevels emerge because the atom’s magnetic moment $\boldsymbol{\mu}$ couples to an external magnetic field \mathbf{B} with $\boldsymbol{\mu} \cdot \mathbf{B} = \mu_z B_z$. The magnetic moment of the ground state μ_z comprises a spin component $\mu_S = g_S m_S \mu_B$ and a nuclear component $\mu_I = g_I m_I \mu_N$, where $g_{S,I}$ denotes the g -Factor, μ_B is the Bohr magneton and μ_N is the nuclear magneton [8, 136, 137]. m_S

*Toptica ECDL DL pro

†8 W, 1178 nm Raman Fiber Amplifier by MPB Communications Inc.

and m_I denote the magnetic quantum numbers of the spin \hat{S}_z and nuclear-spin \hat{I}_z projection along \mathbf{B} , respectively. To solve for the magnetic energy eigenstates across all fields, the full Hamiltonian \hat{H} of the hyperfine structure must be solved, which reads [8, 103, 141]

$$\hat{H} = A \frac{\hat{\mathbf{I}} \cdot \hat{\mathbf{S}}}{\hbar^2} + \frac{1}{\hbar} (g_S \mu_B \hat{\mathbf{S}} - g_I \mu_N \hat{\mathbf{I}}) \cdot \mathbf{B}, \quad (3.1)$$

where A is the hyperfine constant, dependent on the fine structure level [136, 137]. The solution to the hyperfine Hamiltonian in an external magnetic field is shown in Fig. 3.2 for the ground states of both isotopes ${}^6\text{Li}$ (left) and ${}^{23}\text{Na}$ (right).

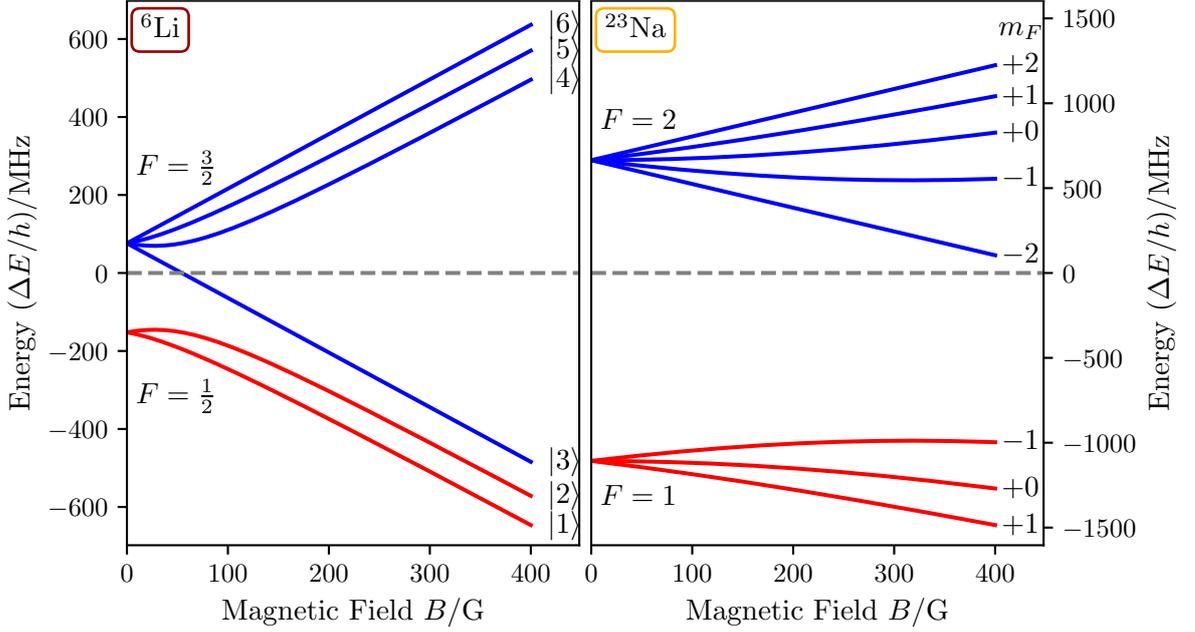


Figure 3.2: Breit-Rabi diagrams for the $2^2\text{S}_{1/2}$ ground state in ${}^6\text{Li}$ and the $3^2\text{S}_{1/2}$ ground state in ${}^{23}\text{Na}$. The figures depict the energy shift ΔE relative to the unperturbed fine structure ground state (dashed line) with respect to the external magnetic field B .

In the limit of very small magnetic fields $A \gg g_S \mu_B B$, the splitting is centred around the $F = 1/2$ ($F = 1$) and $F = 3/2$ ($F = 2$) hyperfine states of lithium (sodium). Here, both the total angular momentum F and its projection m_F serve as good quantum numbers, with each F possessing $2F + 1$ distinct m_F states of energy $g_F m_F \mu_B B$, visible in the initial linear slopes. However, the eigenstates m_F are mixtures of the original basis states, obeying $m_F = m_S + m_I$. As the magnetic field increases, this leads to strong bending of states that do not satisfy $|m_F| = I + S$ (maximally stretched states).

In the high magnetic field limit $A \ll g_S \mu_B B$, the spin's magnetic moment μ_S and the nucleus's magnetic moment μ_I precess separately around the external magnetic field \mathbf{B} , requiring m_S and m_I as labels once again. In this regime, the states split into a high-field-seeking manifold of $m_S = -1/2$ and a low-field-seeking manifold of $m_S = 1/2$. Nonetheless, the notation $|F, m_F\rangle$ still adequately describes the state $m_F = m_I + m_S$.

Throughout this thesis, we adopt the conventional labelling for lithium, designating the

states from lowest to highest energy as $|1\rangle$ to $|6\rangle$, as depicted in Fig. 3.2. For sodium, we stick to labeling the states with their respective $|F, m_F\rangle$ quantum numbers.

3.2 Preparation of an ultracold fermionic sample

In this section, we describe the apparatus that makes up the experiment, and that provides us with an experimental cycle of roughly 26s from hot atoms to a degenerate spin mixture (including cooling of coil elements afterwards). The apparatus is centred around a vacuum chamber, which is essential for preparing dilute and ultracold samples with reasonable lifetimes of a few seconds. Thus, we begin by describing the vacuum system. Since the alkalis used in this experiment come as solid metal chunks at room temperature, the preparation involves multiple steps to create the sample, which are covered subsequently.

3.2.1 Vacuum system

A 3D render of the vacuum chamber forming the experimental apparatus is shown in Fig. 3.3. Atoms start in their respective reservoirs on the left-hand side, labelled as ^{23}Na and ^6Li . This part of the chamber, called the “oven chamber”, heats up the atoms to create a mixed vapour. Next, the mixed isotope gas of hot atoms travels through the Zeeman slower, where it is decelerated to be captured in the magneto-optical trap (MOT) within the “main science chamber”. Here, several cooling and trapping stages follow, ultimately producing a degenerate spin mixture in the main science chamber, spatially offset from the MOT position. The preparation process is detailed in the subsequent sections of this chapter.

Four ion getter pumps[†] are distributed along the system to maintain a vacuum. Additional titanium-sublimation pumps are installed along the setup, which are optionally fired when the vacuum starts to deteriorate. Along the Zeeman slower, a differential pumping stage allows for much better vacuum in the main science chamber, where low pressure is most critical [40]. Due to the high vapour pressure in the oven chamber, the pressure here is around 5.5×10^{-8} mBar. After the differential pumping stage, pressures below 1×10^{-10} mBar are achieved in the main science chamber.

Now, we examine the path of the atoms through the adjoining segments of the vacuum chamber in greater detail.

3.2.2 Effusive Oven

The first stage of the vacuum system is commonly referred to as the “oven chamber”. In our experiment, we use an isotope mixture of ^6Li and ^{23}Na , which necessitates a multi-species effusive oven with individual temperature control for each isotope. The oven chamber design is inspired by the setup built at MIT [142] for the same mixture. In my master’s thesis, I have designed and built an updated version that reduces the maintenance frequency and includes a shutter for the atomic beam [143]. This updated version has been installed during the second-to-last sodium exchange in the summer of 2022 and is shown in the 3D render in Fig. 3.3.

[†]Gamma Vacuum

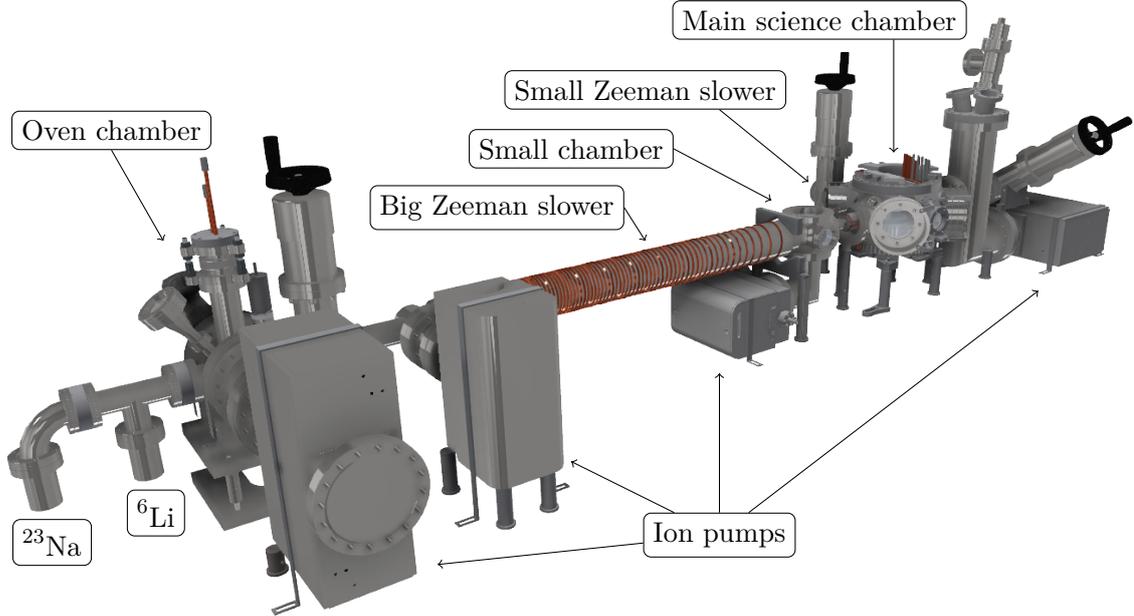


Figure 3.3: Render of the vacuum system. The atoms start on the left-hand side of the setup in a mixed oven configuration, forming an atomic beam. The beam is then slowed down in the spin-flip Zeeman slower. Finally, a mixture of ^{23}Na & ^6Li is trapped and further cooled to quantum degeneracy in the main science chamber on the right-hand side. Across the entire apparatus, four ion pumps maintain a vacuum. Render created with *Autodesk Inventor*.

The oven comprises two individual reservoir cups for ^6Li and ^{23}Na . Each cup is heated well above the respective melting points of 97.80°C for ^{23}Na and 180.54°C for ^6Li to achieve high vapour pressures. During operation, we set the temperature of the sodium reservoir to 330°C and the lithium reservoir to 400°C . In standby mode, temperatures are reduced to 240°C and 300°C , respectively, minimising atom consumption.

The reservoirs are filled with $\sim 10\text{ g}$ of enriched[§] ^6Li and roughly 50 g of[¶] ^{23}Na . After extended operation, the reservoirs eventually run empty. For example, ^{23}Na typically lasts about 2.5 y with a 50 g supply. To refill, we close the gate valve between the oven chamber and the Zeeman slower, flood the chamber with argon, and open the reservoirs. While ^6Li has not yet run empty, inspection during the 2022 instalment of the upgraded oven chamber (with a new lithium supply) have revealed it was nearly depleted, indicating a turnaround time of $\sim 8\text{ y}$. The most recent sodium refill occurred in January 2024, when 50 g of new sodium was inserted.

A narrow nozzle between the reservoirs regulates relative vapour pressures and prevents backflow to the sodium cup. In the big oven chamber, the exit nozzle after the lithium reservoir and a water-cooled copper aperture help to collimate the atomic beam. Atoms that diverge from the intended trajectory adhere to the copper plate, reducing background pressure and stabilising the atomic beam. The beam is then entering the Zeeman slower.

[§]Sigma Aldrich - 340421 - 95 % enriched

[¶]Sigma Aldrich - 282065

3.2.3 Zeeman slower

The Zeeman slower is the first stage for cooling the mixed atomic beam emerging from the effusive oven chamber. In general, the idea of the Zeeman slower is to keep the atoms on resonance with a closed-cycle transition of a laser beam counterpropagating to the atomic beam. Atoms can either populate the excited or ground state of this transition. When in the ground state, an atom absorbs a counterpropagating photon, resulting in a momentum change of $\hbar k$ opposing its initial direction, where k is the wave vector of the laser field. Atoms in the excited state can undergo spontaneous or stimulated emission. Since spontaneous emission lacks a preferred direction, its momentum change averages to zero. The total force F exerted on the atoms is given by the momentum change $\hbar k$ multiplied by the absorption rate, resulting in

$$F = \hbar k \frac{\Gamma}{2} \frac{I/I_{\text{sat}}}{1 + I/I_{\text{sat}} + 4\delta^2/\Gamma^2}, \quad (3.2)$$

where I/I_{sat} represents the intensity in terms of the isotope's saturation intensity [136, 137], Γ denotes the excited-state decay rate and δ is the detuning from the cooling transition. For very high intensities I/I_{sat} , the populations of the ground and excited states are in equilibrium and the force takes its maximum value of $F = \hbar k \Gamma / 2$.

As the atoms slow down, the (blue) Doppler shift of the laser frequency as seen by the atoms reduces, moving the atoms out of resonance. To address this, the Zeeman slower has a magnetic field B that varies along its slowing axis, synchronising the Doppler shift with the change in resonance frequency due to the Zeeman effect. The total detuning of both effects is given by

$$\delta = \delta_{\text{lab}} + kv - \mu_{\text{eff}} B / \hbar, \quad (3.3)$$

where v is the atom's velocity and μ_{eff} denotes the difference in magnetic moments between the excited and ground state.

Our experiment employs a spin-flip Zeeman slower with a variable pitch helix coil design, as described in [144]. In this design, the magnetic field variation is realised through the coil's changing pitch, with the highest field near the oven chamber, a zero crossing in the "small chamber" (refer to Fig. 3.3), and a smaller field of opposite sign at the exit of the Zeeman slower, facing the main science chamber. The spin-flip Zeeman slower comprises two separate segments called the "big" and "small" Zeeman slower, as seen in Fig. 3.3. This design allows for lower peak magnetic field amplitudes, and since the finite magnetic field at the end of the Zeeman slower ends abruptly, the slowed atoms quickly leave the resonance condition and become transparent to the cooling light field. Red-detuned cooling light is coupled through a window at the end of the chamber and travels through the entire Zeeman slower. The closed transitions are driven by σ^+ polarised light between the maximally stretched states, as illustrated in Fig. 3.1. For ${}^6\text{Li}$, this is the $2^2\text{S}_{1/2} |F = 3/2\rangle \rightarrow 2^2\text{P}_{3/2} |F = 5/2\rangle$ transition, and for ${}^{23}\text{Na}$, it is the $3^2\text{S}_{1/2} |F = 2\rangle \rightarrow 3^2\text{P}_{3/2} |F = 3\rangle$ transition. Though the closed transition should prevent substantial occupation of dark states, a small admixture of repumper light is added to enhance the efficiency.

After passing through the Zeeman slower, the atoms are sufficiently cooled for loading into the magneto-optical trap.

3.2.4 Magneto-optical trap

After the Zeeman slower, the atoms are loaded into the magneto-optical trap (MOT). Both isotopes are loaded simultaneously in overlapping but slightly displaced MOTs.

The working principle of a MOT relies on the same force as the Zeeman slower, introduced in Eq. (3.2). Here, we combine the effect of a molasses, which counteracts the atom's velocity to slow them down, with a spatially restoring force that creates a trapping potential [103, 141]. This is achieved with six red-detuned beams for each isotope, consisting of two counterpropagating beams along each Cartesian coordinate. The counterpropagating beams create a molasses force proportional to the atom's velocity. Atoms moving towards a beam experience a blueshift in that beam, resulting in a higher scattering rate and net force (Eq. (3.3)). Combining this effect across all three Cartesian axes results in a 3D molasses.

The spatial trapping potential is realised by adding a magnetic field and selectively driving closed transitions with polarised light. To this end, a quadrupole magnetic field $B(x, y, z) \propto (x, y, -2z)^\top$ is created using a coil pair in an anti-Helmholtz configuration. Due to the Zeeman splitting of the different magnetic sublevels (see Fig. 3.2), the MOT light is shifted into resonance for atoms moving out from the trap centre towards higher magnetic fields. This results in a spatially restoring (trapping) force towards the trap centre, proportional to the displacement.

The MOT centre positions are in line of sight of the atomic beam and are defined to first order by the zero magnetic field centre of the MOT coils. Small displacement between the MOTs can be introduced by adjusting the power balance of counterpropagating MOT beams. Later, the zero-field position defines the start of the small magnetic trap before the minimum position is moved to the big magnetic trap, and so large displacements can reduce the transport efficiency significantly. However, small displacements through power balancing can result in a reduction of light-assisted collisions [145], increasing the atom number in both MOTs.

For ${}^6\text{Li}$, our MOT closely resembles the textbook example discussed in [141]. Here, a small amount of repumper light is added to the cooling light to reduce the population of dark-states. In contrast, the ${}^{23}\text{Na}$ MOT employs a technique called the “dark spot MOT” [146], in which repumper light is only admixed in a ring centred around the MOT (creating a dark spot). This allows atoms in the centre to have a macroscopic population of dark states, reducing light assisted collisions and thereby significantly increasing densities in the trap centre.

Molasses and compressed MOT

At the end of the MOT phase, after shutting down the Zeeman slower and fully loading the MOTs (around $\sim 2.2\text{s}$), the ${}^6\text{Li}$ MOT is compressed to increase the phase-space density, and therefore, decrease the temperature. This is achieved by ramping down the detuning and power of the lithium MOT beams while keeping the quadrupole gradient constant, a process known as “compressed MOT” (cMOT).

Subsequently, the quadrupole field, the lithium light as well as the dark spot light for sodium are turned off. Simultaneously, sodium repumper light is added to the sodium cooling light, realising a pure optical molasses for ${}^{23}\text{Na}$. This technique works via the principle of polarisation gradient cooling, as described in [147, 148], and is used for around 2 ms. The

polarisation gradient forms due to interference of the counterpropagating σ^\pm cooling beams.

The theoretical minimum temperature in the MOT is limited by the Doppler temperature $k_B T_D = \hbar\Gamma/2$ [141], which originates in random walk events due to spontaneous emissions, as discussed in Section 3.2.3. For the considered cooling transitions, this corresponds to $T_D = 140 \mu\text{K}$ for ${}^6\text{Li}$ [136] and $T_D = 240 \mu\text{K}$ for ${}^{23}\text{Na}$ [137]. However, polarisation gradient cooling can achieve temperatures below this limit, reaching the recoil limit [141, 147]. After the cMOT and molasses stages, the temperatures of both isotopes have been estimated by time-of-flight measurements to be $290(30) \mu\text{K}$ (${}^6\text{Li}$) and $125(8) \mu\text{K}$ (${}^{23}\text{Na}$) [40], containing approximately 10^9 and 10^{10} atoms, respectively.

In the next stage, the mixed gas is transferred to the magnetic trap for further cooling. For this, optical pumping is performed to prepare atoms in the maximally stretched and low-field-seeking states $|6\rangle$ for ${}^6\text{Li}$ and $|F, m_F\rangle = |2, 2\rangle$ for ${}^{23}\text{Na}$ (see Fig. 3.2). The optical pumping begins as soon as the molasses stage is completed, with all cooling light and the MOT quadrupole field turned off. A small bias field is added to preserve a quantisation axis.

3.2.5 Magnetic trap

The magnetic trap is the next step for further cooling, exploiting the lifted degeneracy of the magnetic sublevels in an external magnetic field, as shown in Fig. 3.2. In order to trap atoms in a magnetic quadrupole field $B(x, y, z) = B'(x, y, -2z)^\top$, they must occupy low-field-seeking states, which is achieved by optical pumping to the maximally stretched and low-field-seeking states $|6\rangle$ in ${}^6\text{Li}$ and $|F, m_F\rangle = |2, 2\rangle$ in ${}^{23}\text{Na}$ at the end of the molasses stage. This results in a linear potential for the states of both isotopes

$$U(x, y, z) = \mu_B B' \sqrt{x^2 + y^2 + 4z^2}, \quad (3.4)$$

where B' is the magnetic gradient in x and y directions, and μ_B denotes the Bohr magneton. The factor $g_F m_F$ equals to $g_S m_S = 1$ for maximally stretched states, and hence drops out. While other state combinations also have low-field-seeking segments for small magnetic fields (e.g., sodium [149]), the chosen state combination has proven effective, since spin-exchange collisions to anti-trapped states are forbidden. Additionally, maximally stretched states remain low-field-seeking across all fields.

Initial magnetic trap

The first magnetic trap is realised with the same coils that have been used for the MOT in the previous step. After optical pumping, the magnetic gradient is turned back on, optimised for ideal mode matching between the MOT and magnetic trap, achieving a trap depth of around $\sim 2.5 \text{ mK}$ [40].

At this stage, the magnetic trap centre is still overlapped with the initial MOT, but this only presents the first stage of the magnetic trap. Next, the atoms are moved from the initial magnetic trap to the big magnetic trap created with the ‘‘Feshbach coils’’ in quadrupole configuration, later used in Helmholtz configuration to produce large, uniform magnetic fields of up to 1000 G at the atoms’ position to employ Feshbach resonances, as described in Section 2.2.2. The centre position of the final magnetic trap is just 3 mm below the upper

viewport of the main science chamber and allows for future imaging with a high numerical aperture or the superposition of a rapidly changing magnetic field using a small, specialised coil, which will be described in Section 4.5.

For transport to the final magnetic trap, the gradient of the MOT coils is slowly turned off while the current of the Feshbach coils is being ramped up, moving the atoms by more than 30 mm. The precise functions of the MOT and Feshbach currents for this transport have been optimised in [41].

Final magnetic trap & radio-frequency induced evaporation

In the final magnetic trap, evaporative cooling is achieved via radio-frequency (RF) induced spin-slips in ^{23}Na . For a reasonable evaporation efficiency, the trap is turned on for more than 5 s, and thus additional loss and heating channels must be addressed and mitigated.

In the centre of the quadrupole trap, the magnetic field linearly approaches zero, allowing spin-flip transitions to anti-trapped high-field-seeking states to occur – a phenomenon known as Majorana losses [150]. Since cold atoms spend more time near zero potential in the trap centre, cold atoms are more prone to be lost, resulting in heating [151, 152]. To mitigate this, the final magnetic trap combines the magnetic quadrupole field (Eq. (3.4)) with a blue-detuned optical “plug” potential, which repels atoms from the trap centre [153]. The plug laser^{||} runs at 532 nm with a maximum power of 15 W. Due to degradation, the laser output is currently limited to about > 10 W, which still seems to not limit the final atom number in the trap.

Evaporation is performed by inducing RF transitions between initial low-field-seeking and final high-field-seeking states [30, 86]. Atoms transferred to high-field-seeking states are expelled from the trap, allowing the remaining gas to thermalise anew. If atoms with energies above the sample’s average energy are removed, the temperature characterising the energy distribution decreases (see distributions in Section 2.1). Now, if the amount of lost atoms drops slower than gain in the thermal de-Broglie wavelength, an increase in phase-space density is achieved [85]. This principle of evaporative cooling is explored further in Chapter 7, where it is implemented in a novel way to a gas that has already achieved degeneracy.

Here, evaporation involves a transition between the initial $|F = 2, m_F = 2\rangle$ state and the anti-trapped, maximally stretched and high-field-seeking state $|F = 1, m_F = 1\rangle$ (Fig. 3.2). In order to address the transition for arbitrary magnetic fields B , a frequency of $1771.6 \text{ MHz} + B \cdot 1.4 \text{ MHz/G}$ must be applied (refer Fig. 3.2). For evaporative cooling, high-energy atoms near the edge of the trap (largest magnetic fields) must be removed to lower the temperature. As the temperature decreases, the average energy reduces, and high-energy atoms are found closer to the trap centre. Accordingly, the RF frequency is linearly ramped from 1900 MHz to 1772.5 MHz over a span of 5 s [40, 41]. The optimal evaporation parameters have been determined phenomenologically by optimising for the highest atom number in the final degenerate sample. Here, initial evaporation efficiencies of $\gamma = 1.3$ and $\gamma = 3.7$ for ^{23}Na have been measured [41].

After evaporation, no sodium remains, while the polarised $|6\rangle$ state ^6Li cloud has been

^{||}Lighthouse Photonics - Sprout-G-15W

sympathetically cooled with minimal losses (sympathetic evaporation efficiency of $\gamma = 4.9$ [41]). The post-evaporation temperature of ${}^6\text{Li}$ has been determined to be below < 20 uK [40]. Next, the polarised ${}^6\text{Li}$ sample is loaded into the optical dipole trap.

3.2.6 Dipole trap

The dipole trap presents the final stage in the preparation of a degenerate two-state mixture of ${}^6\text{Li}$. By this stage, all sodium has been consumed during the sympathetic evaporative cooling of ${}^6\text{Li}$ in the magnetic trap. Now, all ${}^6\text{Li}$ is transferred into the optical dipole trap.

Dipole traps work on the principle of the AC-stark shift. A light field of frequency ω detuned by $\delta = \omega - \omega_0$ from the transition frequency ω_0 imposes a light shift on the transition levels [154]. The ground state energy is shifted downwards for negative (red) detuning $\delta < 0$. This creates an attractive potential $U(\mathbf{r})$ for the atoms, proportional to the intensity $I(\mathbf{r})$ of the laser field, which reads [155] (in the limit of very large detuning $|\delta| = |\omega - \omega_0| \gg \Omega_R$, with respect to the Rabi-frequency [103])

$$U(\mathbf{r}) = \frac{\hbar\Gamma^2}{8\delta} \frac{I(\mathbf{r})}{I_{\text{sat}}}, \quad (3.5)$$

where Γ is the decay rate of the excited state and $I(\mathbf{r})/I_{\text{sat}}$ denotes the laser intensity with respect to the transition dependent saturation intensity I_{sat} [136].

As indicated by the spatial dependence of the potential $U(\mathbf{r})$ and proportional intensity $I(\mathbf{r})$, the potential takes the shape of the underlying laser field. This is an extremely powerful feature, as potentials can now freely be designed by laser beam shaping, which has been done for various geometries, e.g., homogeneous box traps [156, 157], optical tweezer arrays [158] or lattice geometries [159].

Here, we use a crossed optical dipole trap composed of two intersecting TEM_{00} Gaussian beams with a respective intensity profile of [103, 160]

$$I(x, y, z) = I_0 \exp\left(-\frac{2x^2}{w_x^2} - \frac{2y^2}{w_y^2}\right). \quad (3.6)$$

The z (propagation direction) dependency has been omitted, as the intensity profile along z changes on the order of the Rayleigh range $z_R \approx 40$ mm [40], while the cloud's extent is limited to within 3 mm by the chamber windows. Typically, clouds below 1 mm in size are prepared. Here, $I_0 = 2P/(\pi w_x w_y)$ represents the centre intensity of an elliptical beam with power P , and waists w_x and w_y in x and y direction, respectively [160].

Plugging Eq. (3.6) into Eq. (3.5) yields the full potential of the Gaussian beam. In this experiment, the optical dipole potential is formed by two intersecting Gaussian beams – one in the horizontal plane and the other vertically orientated [40, 41]. Close to the atoms' position, the full potential can be simplified via a harmonic approximation, which brings the potential to the known form from Eq. (2.27) and reads [103]

$$U(x, y, z) = \frac{1}{2}m\left(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2\right), \quad (3.7)$$

where $m = m_{\text{Li}}$ is the mass of the trapped ${}^6\text{Li}$ atoms. Interestingly, the squared trapping frequencies are proportional to the beam power $\omega_i^2 \propto I$ [103]. However, The full Gaussian potential remains important for the wings of the trap, where the potential deviates from the harmonic approximation and thermometry measurements are conducted in Chapter 5.

The dipole trap is created by a red-detuned ytterbium fibre laser** operating at 1070 nm with a maximum output power of 50 W. Light from this laser is coupled via two individual fibres for the horizontal and vertical beam [40, 41]. An AOM is placed in front of each fibre for power stabilisation and to run ramps of the trapping potential. Details about the setup can be found in the construction thesis [161].

Preparation of a spin mixture

Immediately after evaporation in the magnetic trap (as discussed in Section 3.2.5), the dipole trap powers are ramped up, and the magnetic trap is ramped down. ${}^6\text{Li}$ in the $|6\rangle$ state now occupies the dipole trap, and magnetic fields can be applied freely using the big Feshbach coils, producing magnetic fields of up to 1000 G. To create uniform magnetic fields, the Feshbach coils are switched to a Helmholtz configuration from the quadrupole configuration used during the magnetic trap stage, employing IGBTs in an H-bridge configuration [40, 41].

To further cool the ${}^6\text{Li}$ sample, a two-state mixture is now required, as the Pauli exclusion principle prevents scattering in a polarised Fermi gas. To this end, the $|6\rangle$ state is first transferred to the $|1\rangle$ state via a Landau-Zeener sweep [8, 162]. A second Landau-Zeener sweep then creates a balanced and coherent superposition state $(|1\rangle + |2\rangle) / \sqrt{2}$. This coherent state is also protected from scattering by the Pauli principle, so a short decoherence time is introduced to end up with an equal mixture of $|1\rangle$ and $|2\rangle$, which we denote as a $|12\rangle$ mix [40, 41]. Decoherence is introduced by slight magnetic field inhomogeneities, which mixes accumulated phases of the atoms and destroys the coherence after ~ 100 ms [40]. In some experiments, a $|13\rangle$ mixture is preferred due to a narrower (still broad) Feshbach resonance (Fig. 2.3), which is prepared by adding an additional transfer from $|2\rangle$ to $|3\rangle$ before the decoherence time.

Feshbach evaporation

With a balanced spin mixture, the Pauli exclusion principle no longer prohibits interactions in the sample. Interactions are tuned via the uniform magnetic field produced by the Feshbach coils in Helmholtz configuration.

For fast and efficient thermalisation, the magnetic field is ramped close to the Feshbach resonance of the prepared sample. Here, the power of the optical dipole trap is exponentially ramped down, preferably removing high-energy atoms.

This process yields a degenerate $|12\rangle$ or $|13\rangle$ spin mixture of up to $N_\sigma = 1 \times 10^6$ atoms per spin state at a temperature of roughly $0.07T_F$ [40, 41]. The trap frequencies are typically around $2\pi \cdot 100$ Hz and $2\pi \cdot 300$ Hz, measured by exciting dipole modes of the cloud's centre-of-mass [8, 103].

**IPG Photonics, YLR-50-LP

3.2.7 Quantum simulator

From here on out, the desired scattering length can be set by adiabatically adjusting the external magnetic field. Values for the scattering length are calculated via the data from [111], shown in Fig. 2.3, using the calibrated magnetic field at the atoms' position.

Additional compensation coils allow for the application of magnetic gradients and magnetic compensation fields, e.g., to compensate for gravity. Additionally, the trap can be adjusted by setting ramps for the dipole beam powers. This allows for probing the BEC-BCS crossover and is explored in the later chapters.

Magnetic field calibration

For a precise knowledge of the magnetic field and the scattering length, the magnetic field requires precise calibration. A rough calibration has been performed via measuring the field before assembly of the experiment via a Hall probe [40, 41]. Now, a precise calibration is performed through the exact knowledge of the $|1\rangle$ to $|2\rangle$ transition frequency, as shown in the Breit-Rabi spectrum in Fig. 3.2.

To this end, a polarised Fermi gas in the $|1\rangle$ state is prepared. Next, the transition frequency $\nu_{1\rightarrow 2}$ to the $|2\rangle$ state is probed at various magnetic fields B , which is indicated by strong losses of the $|1\rangle$ state in absorption images (refer to Section 4.1). The resonance position is then fitted to the known spectrum, applying a second-order^{††} calibration function of the magnetic field, as shown in the left panel of Fig. 3.4. On the right side of the figure, the

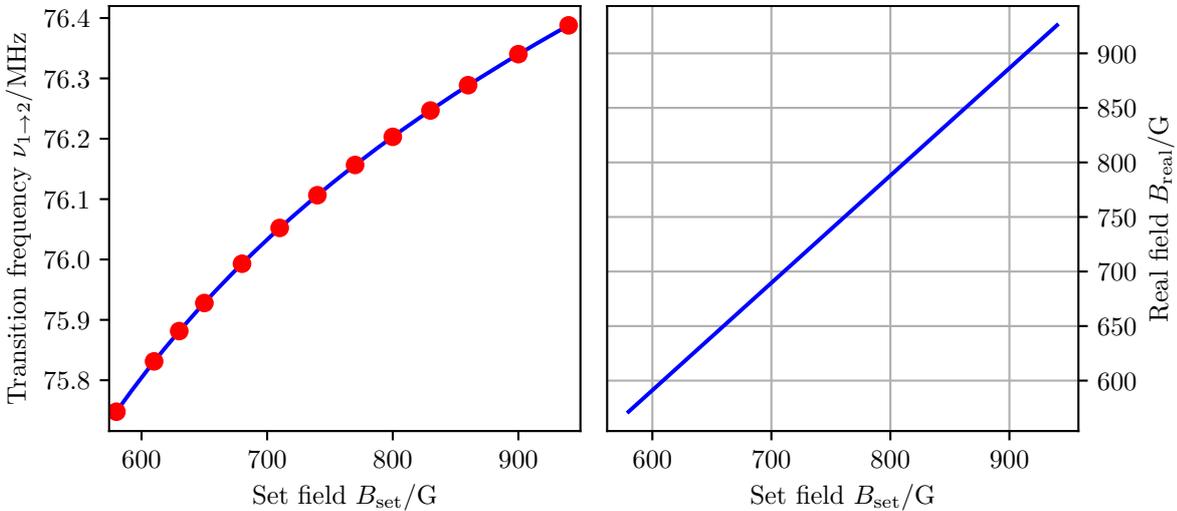


Figure 3.4: Calibration of the Feshbach field. The calibration has been performed by measuring the RF resonance position of the polarised $|1\rangle$ to $|2\rangle$ transition for different (magnetic) set fields B_{set} . At resonance, a strong loss in the absorption image of the $|1\rangle$ state has been observed. By fitting (solid line, **left**) the determined resonance frequencies $\nu_{1\rightarrow 2}$ (red dots, **left**) to the expected Breit-Rabi spectrum in Fig. 3.2, a second-order polynomial calibration function (Eq. (3.8)) of the magnetic field has been determined (**right**). Errors of the transition frequency are too small to be visible.

^{††}Inclusion of a third-order changes the calibration function on the order of 1×10^{-5} and is not resolvable considering the magnetic field noise

fitted calibration function is presented, providing a mapping of the set value of the magnetic field B_{set} (from rough calibration) to the real magnetic field B_{real} at the atoms' position. The fitted calibration function yields

$$B_{\text{real}}(B_{\text{set}}) = 0.98563(6) \cdot B_{\text{set}} - \frac{1}{2} \cdot 1.7(2) \times 10^{-6} \text{ G}^{-1} \cdot B_{\text{set}}^2. \quad (3.8)$$

The stated error (0.01 %) represents the standard error of the fit. To diagnose the magnetic field value throughout the experimental cycle, the current passing the coils is measured with a high-precision current transducer^{††}, yielding a voltage proportional to the current. Using this probe, we find the noise of the magnetic field to be matching the calculated error of 0.01 %. Subsequently, the calibrated value is used for precise determinations of the scattering length and bound state energy, which becomes important later in Chapter 7.

Dipole power calibration

The power of the dipole beams is stabilised by an AOM in front of each out-coupling fibre. Here, the RF power going to the AOMs is regulated via a photodiode at the end of each fibre with respect to a set point, which is changed to desired values throughout the experimental cycle, e.g., to drive ramps. To obtain a rough idea of the power in each beam as a function of the set point voltage, a power calibration has been performed. To this end, the power in each beam has been measured with a power meter^{§§} as a function of the regulation voltage. Linear fits of the calibration measurements are shown in Fig. 3.5.

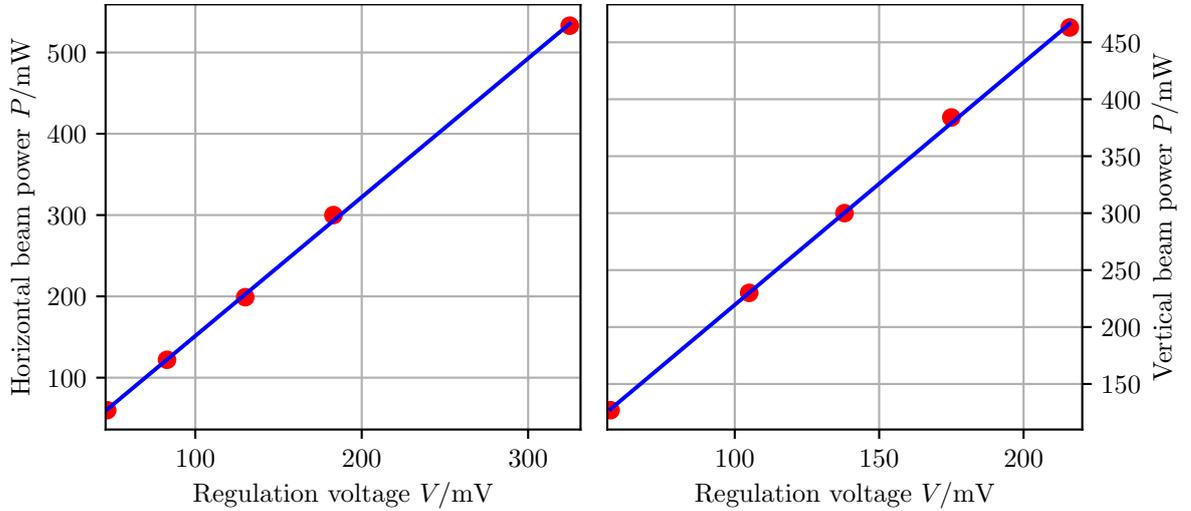


Figure 3.5: Calibration of the dipole beam powers. The power calibration for the horizontal (**left**) and vertical (**right**) beams as a function of the photodiode regulation (set) voltage is shown with the raw data (dots) and corresponding linear fits (solid line).

^{††}LEM, ITN 600-S ULTRASTAB

^{§§}Ophir, PD300-3W

Detection and Methodology

This chapter presents the techniques used to detect, quantify and excite the Fermi gas phases studied in this thesis. As these are essential aspects of the quantum simulator, most techniques have been previously discussed and analysed in prior works on this experiment [133, 135, 163]. Here, the most important aspects are revisited and discussed, necessary to perform and quantify the experiments in Chapters 5 to 7.

Detection is performed via resonant absorption imaging of the quantum gas after preparation and optional excitation. Resonant imaging of the gas creates a shadow in the imaging plane, which corresponds to the optical density of the gas. Since the density of the gas is proportional to the optical density, the optical density must be determined accurately, requiring a thorough calibration of the imaging system.

In the next step, the in-situ density – measured without thermal expansion – is used to perform thermometry of the strongly interacting Fermi gas. From this, we acquire thermometric values of the gas such as the 3D density, temperature and chemical potential. Owing to the geometry of imaging from a single direction, the absorption image can only provide the density integrated along the imaging direction. However, because of the axial symmetry of the trapping potential, the inverse Abel transformation [164] allows to reconstruct the 3D density of the inhomogeneous gas. By knowledge of the equation of state (EoS) at unitarity at the superfluid critical point [165], the 3D density can be employed to reconstruct the underlying trapping potential. With the virial expansion of the equation of state providing an approximate solution for the low-density regions of the cloud, temperatures and chemical potentials can then be determined throughout the entire crossover.

While the determination of the thermometric parameters of the gas is a key aspect to quantify the performed experiments, it does not give a clear signature of crossing the superfluid phase transition. Throughout the crossover, the superfluid phase is linked to the existence of a condensate; thus, pairing statistics perturb the momentum distribution in comparison to the non-condensed gas in the normal phase (refer to Section 2.3.2). In the BEC limit of the crossover, thermal expansion indeed reveals a bimodal distribution of the condensed zero-momentum dimers and the remaining thermal dimers [38, 49]. However, in a strongly interacting Fermi gas, the absorption image of the expanding gas lacks a clear signature of superfluidity. This is because the Fermi edge smears out due to competing effects such as temperature, interactions and the inhomogeneity of the trap [8]. To address this, a

state-of-the-art technique known as the rapid ramp (RR) of the magnetic field is introduced. This technique constitutes a projection of weakly bound pairs onto tightly bound dimers, and restores the bimodal distribution throughout the entire crossover [38, 49].

While the rapid ramp has been successfully used to detect condensation across the crossover, its perturbative effect on the condensate has been matter of debate [50–52, 166]. It has been found that the exponential ramp of the magnetic field leads to a non-monotonic dependency of the condensate fraction with the initial magnetic field [50–52]. While the formation of a condensate during the ramp could be ruled out [166], the perturbation could lead to distortions near the onset of condensation. Therefore, alternative approaches to identify a finite condensate have been explored in this thesis. These rely on neural networks for advanced image recognition – to infer the condensate fraction directly from time-of-flight images. To this end, this chapter introduces the basic concepts of the neural networks employed in Chapter 5.

Finally, the device to excite the Fermi gas is explored. In [163], a custom-built small magnetic field coil with a finite magnetic field and a negligible magnetic field gradient at the atoms' position has been presented. This coil is used to superimpose a rapidly changing magnetic field, which allows for changing the scattering length, and thus the interaction parameter of the system. Two modes of operation, quenching and modulating the magnetic field, are discussed in this chapter.

Outline of this chapter

The sections of this chapter are outlined as follows:

- Section 4.1: First, we introduce the concept of absorption imaging, which serves as our detection method. Here, we also describe its implementation in our experimental setup, listing all available cameras.
- Section 4.2: By imaging atoms in-situ, we can recover the 3D density profile (inverse Abel transformation) and determine the temperature of the gas (fits to the virial expansion of the equation of state). The section details the necessary processing and analysis, including the required reconstruction of the trapping potential.
- Section 4.3: This section introduces the rapid ramp (RR), a state-of-the-art technique for detecting superfluidity throughout the BEC-BCS crossover. Here, we also discuss its limitations.
- Section 4.4: We then provide a broad introduction to the general principals of deep convolutional neural networks, used to detect the onset of superfluidity in Chapter 5.
- Section 4.5: Finally, we discuss the custom-built small magnetic field coil, used to excite the Higgs mode in Chapter 6 and to dissociate composite dimers in Chapter 7.

4.1 Absorption imaging

After conducting the desired experiment, the Fermi gas must be probed to extract information in the form of the density distribution. A common technique for this is resonant absorption

imaging. To perform it, a resonant laser pulse illuminates one of the present hyperfine states in the spin mixture. As the resonant cloud absorbs the laser pulse, the atoms cast a shadow onto the laser, which is then detected by a CCD camera at the end of the imaging system.

4.1.1 Imaging system

The imaging system in the experiment comprises four cameras placed along different axes for various purposes. To describe the spatial arrangement of the cameras, we first establish a coordinate system. A convenient choice of a coordinate system presents the quantisation axis along the Feshbach field, which we define as the z axis. Perpendicular to the quantization axis, the x axis is defined by the travel direction of the horizontal dipole trap beam, while the y direction is perpendicular to the other two directions.

Magnification

An optical telescope between the atoms and the camera defines the focus position and magnification factor for imaging. The magnification has been calibrated by observing the free fall of an atom cloud with cameras positioned perpendicular to the direction of gravity. As remaining magnetic field gradients distort the acceleration, two hyperfine states $|3\rangle$ and $|6\rangle$ with equal but flipped magnetic moments (refer to Fig. 3.2) have been observed here, which cancels the effect of any gradient. Cameras with an observation axis parallel to gravity are calibrated through cross-calibration by observing the same cloud with a different camera perpendicular to gravity.

Andor1 - Andor iXon Ultra 897

Along the vertical z direction, two cameras are used for different expansion times. For in-situ images without any expansion time, the atoms' position is at the focus of the optical dipole trap. To this end, the first camera ("Andor1") is focused to this in-situ position and exposed from below the chamber, capturing images with a magnification of $M = 7.26$ and a pixel size of $w = 16.0 \mu\text{m}$.

This camera is mainly used to measure the in-situ density distribution, from which all thermometric values are deduced. Since this requires a precise determination of the optical density, this camera is calibrated for high-intensity imaging in Section 4.1.2.

Andor0 - Andor iXon Ultra 888

During thermal expansion, the dipole trap is shut down, allowing the atom cloud to expand into the residual trapping potential defined by the curvature of the magnetic field from the Feshbach coils. The trap frequency perpendicular to the z axis is $\omega/(2\pi) = 1/T \sim 16 \text{ Hz}$, while the potential along z is anti-confining. After a quarter-period ($T/4 \sim 15 \text{ ms}$) in this curvature, the absorption image reflects the momentum distribution of the sample [167]. During approximately 15 ms of expansion, the cloud also drops by more than $\sim 1 \text{ mm}$ due to gravity. Thus, the atoms leave the focus position of the first camera, Andor1. For this reason, a second camera ("Andor0") along the vertical direction is installed, which is focused on the atoms' position after $\sim 15 \text{ ms}$ of free fall. The light for this camera enters the chamber

from above, and transverses a telescope with a magnification of $M = 4.32$ after exposing the atoms. The camera has a pixel size of $w = 13.0 \mu\text{m}$.

Andor2 - Andor iXon Ultra 888

Both Andor0 and Andor1 image the atoms along the vertical z axis. The first camera (“Andor2”) in the horizontal plane observes the atoms along the y direction with a magnification of $M = 4.51$ and a pixel size of $w = 13.0 \mu\text{m}$. This camera is primarily used to observe dynamics along the vertical direction, e.g., dipole modes, or to measure the aspect ratio of the density distribution between the z and x axes.

During magnification calibration with atoms in a magnetic quadrupole trap of known aspect ratio of 2 (refer to Section 3.2.5), strong astigmatism along the imaging system of this camera has been observed. Given the true aspect ratio in a quadrupole trap is known to be 2, a correction factor of 1.86 has been determined for calculating the aspect ratio via this imaging system. Later, the aspect ratio from this camera is important for calculating the 3D density via the inverse Abel transformation in Section 4.2.1.

Alta0 - Alta U1

The final component of the imaging system, camera “Alta0”, images atoms along the x direction, with a magnification factor of $M = 1.23$ and a pixel size of $w = 9.0 \mu\text{m}$. Its primary function is to debug the plug position relative to the atom cloud position in the magnetic trap (refer to Section 3.2.5) and to give a (magnification) cross-calibration value between the z and y axes.

4.1.2 High-intensity absorption imaging

In absorption imaging, the atom cloud is illuminated with a resonant laser pulse. Since the atoms are on resonance, they absorb the imaging light and cast a shadow on the laser pulse, which is captured by a camera. This first image, which contains the shadow, is called the “atom image”. Next, a second image is taken using the same imaging parameters but without any atoms present, creating a reference image known as the “light image”.

From the atom and light images, the reduced intensity I_{atom} and the reference value I_{light} is deduced for each pixel. The reduction of the intensity dI/dz in the atom image $I_{\text{atom}} - I_{\text{light}}$ is given by the Beer-Lambert law [141]

$$\frac{dI}{dz} = -n\sigma_0 I, \quad (4.1)$$

where n is the atom density and $\sigma_0 = 3\lambda^2/(2\pi)$ the resonant absorption cross-section for a laser pulse of wavelength λ . The law dI/dz suggests an exponential intensity decay along the imaging direction of z . Hence, the total intensity loss at an image pixel (x, y) is the integrated law along the z direction, which defines the optical density $OD_0 := n_{\text{col}}\sigma_0 = -\log(I_{\text{atom}}/I_{\text{light}})$, where $n_{\text{col}} := \int dz n$ is the column density. For low intensities $I_{\text{light}} \ll I_{\text{sat}}$, and with a perfectly closed imaging cycle transition, this method would suffice to determine the column density n_{col} .

However, in practice, the absorption cross-section $\sigma_{0,\text{eff}}$ may deviate from the ideal value $\sigma_0 = 3\lambda^2/(2\pi)$. Additionally, saturation of the excited state reduces the absorption coefficient at high intensities $I_{\text{light}} \sim I_{\text{sat}}$. Finally, these effects culminate in the modified Beer-Lambert law [168]

$$n_{\text{col}}\sigma_0 = -\alpha \log\left(\frac{I_{\text{atom}}}{I_{\text{light}}}\right) + \frac{I_{\text{light}} - I_{\text{atom}}}{I_{\text{sat}}}, \quad (4.2)$$

where α is a correction factor that accounts for imperfections in the imaging system.

In the experiment, we measure the atom and light images with cameras. Instead of intensity values I , the cameras provide a count number C for each pixel. The amount of counts is proportional to the depleted power P in that pixel, and thus $C \propto I\tau$ is proportional to the intensity I times the pulse length τ . This leads to the experimental form of the modified Beer-Lambert law, which reads

$$n_{\text{col}}\sigma_0 = -\alpha \log\left(\frac{C_{\text{atom}}}{C_{\text{light}}}\right) + \frac{C_{\text{light}} - C_{\text{atom}}}{\dot{C}_{\text{sat}}\tau}, \quad (4.3)$$

where C_{atom} and C_{light} are the camera counts of a certain pixel in the atom and light images, respectively, obtained after imaging with a pulse time τ . Note that both the atom and light images are corrected by a dark image C_{dark} , which is subtracted from the raw count arrays $C_{\text{atom},\text{light}} = C_{\text{atom},\text{light}}^{\text{raw}} - C_{\text{dark}}$. The saturation intensity is recast into \dot{C}_{sat} , representing the count rate when illuminated with the saturation intensity I_{sat} .

CCD calibration

While the saturation intensity I_{sat} is known, its corresponding count rate \dot{C}_{sat} is not, requiring a calibration of the count rate at the saturation intensity. This is achieved by illuminating the camera with a light pulse of known duration and power. To prevent clipping losses, the beam size is reduced with a small aperture. Next, the power is measured* after the aperture and as close to the chamber as possible, resulting in 32.2 μW . Since we are interested in the power at the atoms' position (or the saturation count rate with respect to the atoms), we also measure the power as close as possible after the chamber, yielding 27.0 μW . The average power is then used for at the atoms' position. Finally, the total counts are obtained by integrating over all pixels, corrected by the dark image. The number of total counts, normalised to the pulse power, is shown as a function of the pulse duration τ in Fig. 4.1.

A linear fit to the total counts yields a count rate of $0.356(4) \times 10^6 \mu\text{s}^{-1} \mu\text{W}^{-1}$. Interestingly, the linear fit does not pass through the origin; instead, the data suggests a finite delay or rise time of the imaging pulse, likely caused by the rise time of the imaging AOM. This delay, with a value of $\tau_0 = 73(37)$ ns, is considered in the pulse duration $\tau \rightarrow \tau - \tau_0$ when calculating the column density (refer to Section 4.1.2).

The saturation count rate is calculated by considering the effective pixel area $w_{\text{eff}}^2 = (w/M)^2$ with the magnification $M = 7.26$ and pixel size $w = 16.0 \mu\text{m}$ of Andor1 (refer to Section 4.1), as well as the saturation intensity $I_{\text{sat}} = 2.54 \text{ mW/cm}^2$ of the ${}^6\text{Li}$ D₂-line [136]. This results

*Ophir - PD300

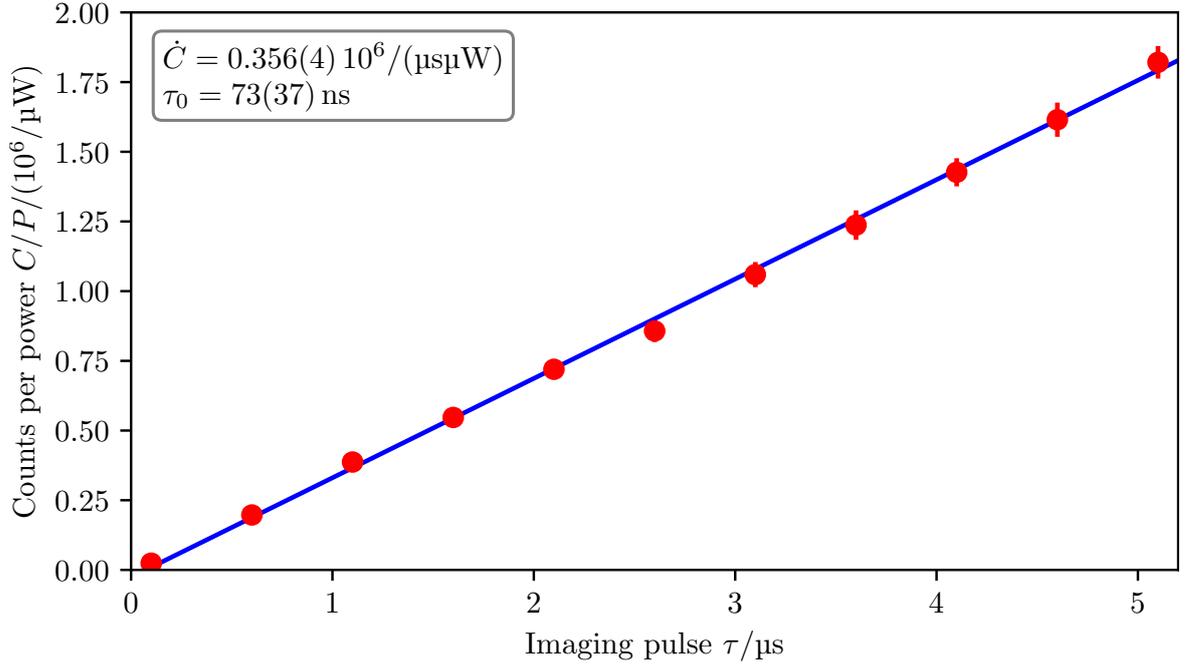


Figure 4.1: CCD calibration of the in-situ camera (Andor1). The plots shows the total counts C , corrected by the dark image and normalised to the power of the imaged beam, as a function of the imaging pulse duration τ .

in a saturation count rate of $\dot{C}_{\text{sat}} = 44.0(5) \mu\text{s}^{-1}$.

Alpha calibration

With the knowledge of \dot{C}_{sat} from the CCD calibration, the modified Beer-Lambert law in Eq. (4.3) can be applied to determine the optical density $OD_0 = n_{\text{col}}\sigma_0$. To this end, the atom cloud is divided into elliptical bins of equal optical density $n_{\text{col}}\sigma_0$ and counts C_{atom} first. Exemplary bins are shown in the upper right corner of Fig. 4.2. Within these bins, the counts from both the atom and light images are averaged to improve the signal-to-noise ratio. Next, the counts are plotted in the form of the modified Beer-Lambert law in Eq. (4.3) by assigning $y := \frac{C_{\text{light}} - C_{\text{atom}}}{\dot{C}_{\text{sat}}\tau}$ the difference term and $x := -\log\left(\frac{C_{\text{atom}}}{C_{\text{light}}}\right)$ the logarithmic term, with $-\alpha$ left as a slope for fitting and the optical density $n_{\text{col}}\sigma_0$ as an offset $y = -\alpha \cdot x + n_{\text{col}}\sigma_0$. An exemplary plot of this is show in Fig. 4.2. To determine both the desired optical density $n_{\text{col}}\sigma_0$ and the correction factor α , each optical density bin is fitted by a linear function with the previously introduced slope ($-\alpha$) and offset ($n_{\text{col}}\sigma_0$).

Finally, both the determined optical density $n_{\text{col}}\sigma_0$ and the correction factor α are presented in Fig. 4.3. The correction factor α shows a linear dependence with the optical density and levels off at $\alpha \sim 2$ for low optical densities $OD_0 < 2$. A similar relationship has already been observed in earlier works in our experiment [40] and in other groups [169].

In principle, the modified Beer-Lambert law in Eq. (4.3) can now be inverted to provide a functional mapping of the counts to the optical density OD_0 [134]. However, we have figured

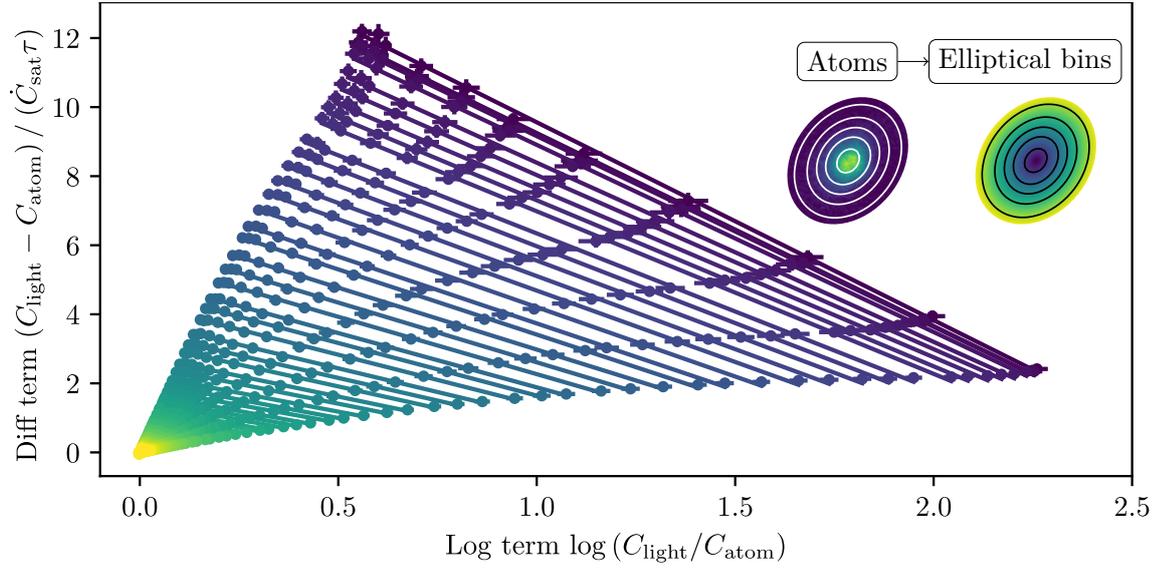


Figure 4.2: High-intensity imaging calibration. The plots shows the application of the modified Beer-Lambert law in Eq. (4.3) by plotting the difference term against the logarithmic term. This leaves $-\alpha$ as the slope and $n_{\text{col}}\sigma_0$ as an offset. This fit requires creating (elliptical) equal (optical-)density bins, shown in the upper right corner. The colour of each linear fit corresponds to a density bin of the same colour.

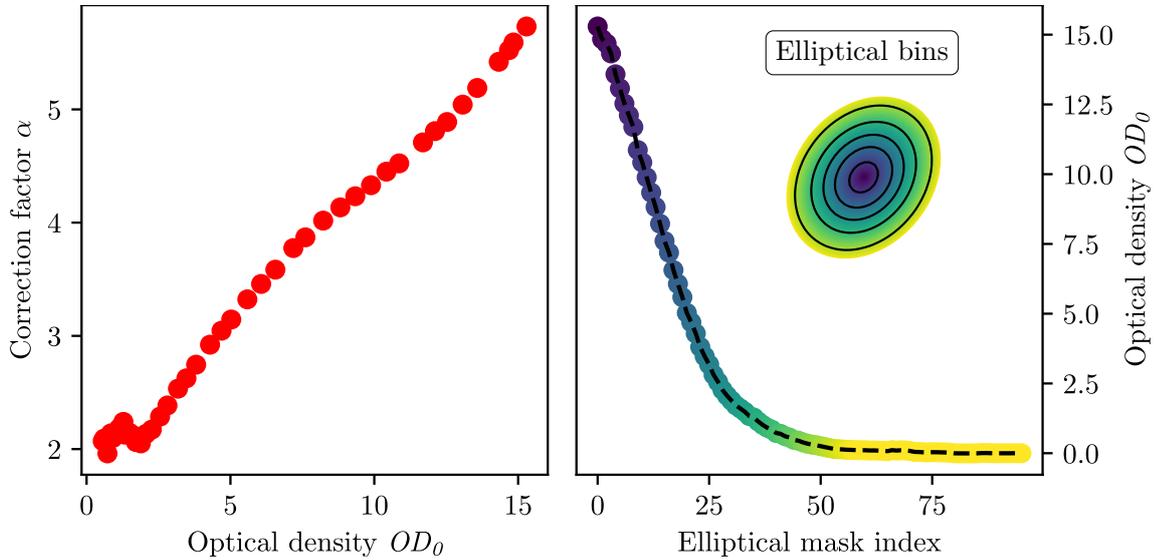


Figure 4.3: α factor and OD_0 from high-intensity imaging. Both panels show the results from the high-intensity imaging calibration shown in Fig. 4.2. **Left:** Dependency of the correction factor α with the optical density OD_0 . **Right:** Optical density OD_0 as a function of the elliptical bin index. The indices are highlighted by the same colour map in the upper right corner. A black dashed line indicates the trend of the optical density.

out that the α calibration depends not only on the optical density OD_0 but also on other experimental parameters, such as the pulse length τ and the magnetic field. The former dependence likely originates from the Doppler effect, while the latter effect can be explained by changes in interactions with the magnetic field. To mitigate this, the α calibration procedure is repeated for each in-situ image, using the fitted offset to determine the optical density, as shown in Fig. 4.2. An exemplary profile of the optical density OD_0 obtained with this technique is shown in the right panel of Fig. 4.3.

Limits of imaging

The previous section considers a stationary atom gas imaged by a resonant laser pulse. However, this scenario usually deviates from reality. As discussed in the context of laser cooling in Section 3.2.3, exposing atoms to resonant light induces a force, as described by Eq. (3.2). This force accelerates and moves the atoms out of the focus position. Additionally, random walk events due to spontaneous re-emission lead to a lateral spread of the gas.

In order to mitigate these effects, the imaging pulse duration τ must be chosen to be short enough. The limit is determined by the combined requirements that (1st) the atoms should not move beyond the depth of focus, (2nd) they should not acquire a Doppler broadened linewidth greater than the power-broadened linewidth and (3rd) the lateral spread should remain below the resolution of the imaging system. For the in-situ camera (Andor1), this poses the strongest limitation. The threshold for this camera is mainly constrained by the relatively shallow depth of focus, approximately $10\ \mu\text{m}$ [133], which results in a maximum pulse duration of $\tau_{\text{max}} \sim 3.5\ \mu\text{s}$ at the highest intensities of $I/I_{\text{sat}} \sim 10$.

4.2 Thermometry of an interacting Fermi gas

In the previous section, high-intensity absorption imaging has been introduced, a technique for determining the column density n_{col} or optical density OD_0 .

This section presents the subsequent steps in image processing to derive the thermometric parameters of the gas. These parameters comprise the temperature T , the chemical potential μ and the spatially-resolved 3D density n_σ per spin state σ of the gas. As discussed in Section 2.4, the density profile of a confined Fermi gas can be described by the local density approximation (LDA), which incorporates the confining potential $U(\mathbf{r})$ into a spatially varying chemical potential $\mu \rightarrow \mu(\mathbf{r}) = \mu - U(\mathbf{r})$. In the LDA, the trap centre can be considered homogeneous, which allows to calculate the (homogeneous) Fermi energy $E_F \propto n_\sigma^{2/3}$ using the density n_σ at the trap centre (refer to Eq. (2.4)). This calculation is particularly important for describing the onset of the superfluid phase transition, as this naturally occurs at the point of the largest Fermi energy first, which is the point of largest density.

For finite temperature T , the density profile of an ideal Fermi gas within LDA is given by the integral in Eq. (2.31). Solving this integral yields [8, 103]

$$n_\sigma(\mathbf{r}) = -\frac{1}{\lambda_{\text{dB}}} \text{Li}_{3/2} \left(-\exp \left(\frac{\mu - U(\mathbf{r})}{k_B T} \right) \right) = -\frac{1}{\lambda_{\text{dB}}} \text{Li}_{3/2}(-z), \quad (4.4)$$

where λ_{dB} denotes the de-Broglie wavelength and $\text{Li}_{3/2}$ (not to be confused with lithium) is

the polylogarithm function of order $3/2$. In the final step, the argument of the polylogarithm function has been replaced by the definition of the fugacity $z := \exp\left(\frac{\mu - U(\mathbf{r})}{k_B T}\right)$ [170].

In order to determine the temperature of the ideal Fermi gas, the straight-forward approach would be to let the gas expand and to measure the extend of the corresponding momentum density distribution via a similar integral to Eq. (2.31). However, this approach is complicated by the introduction of interactions, which are particularly strong in the high-density regions of the gas upon release. Another approach constitutes the in-situ density distribution from Eq. (4.4). In the low-density wings of the gas, the polylogarithm can be approximated by the unity function $\text{Li}_s(x) \approx x$. Here, the density approaches that of a Boltzmann gas (refer to Eq. (2.1)) and the temperature T can be extracted from a fit to the wings via

$$\begin{aligned} n_\sigma(\mathbf{r}) &= \frac{1}{\lambda_{\text{dB}}} \exp\left(\frac{\mu - U(\mathbf{r})}{k_B T}\right) \\ \Leftrightarrow \quad \log(n_\sigma) &= c - \frac{U(\mathbf{r})}{k_B T}, \end{aligned} \quad (4.5)$$

where c is a fitting constant.

Better convergence can be reached by fitting the virial expansion of the equation of state (EoS) to the density profile, which takes interactions into account (refer to Section 4.2.3). Since the fitting slope depends linearly on the trapping potential $U(\mathbf{r})$, a precise determination of the potential is required. This is a challenging process, as the fit requires a precise understanding of the trapping potential in the low-density regions where the harmonic approximation of the potential loses validity (refer to Section 3.2.6).

In this section, we summarize the processing steps from the column density n_{col} to the final thermometric values. This involves reconstructing the 3D density n_σ from the original column density n_{col} , calibrating the trapping potential $U(\mathbf{r})$ and fitting the virial expansion of the EoS to determine the temperature.

4.2.1 Inverse Abel transformation - Centre density

The first step in extracting the thermometric parameters of the gas is the reconstruction of the 3D density n_σ from absorption images. As discussed in Section 4.1, absorption imaging yields the column density $n_{\text{col}} = \int dz n_\sigma$, representing the integral of the 3D density n_σ along the imaging direction. This principle is illustrated in Fig. 4.4.

In order to retrieve the original 3D density, an inverse transformation has to be applied. This transformation is the inverse Abel transformation, which exploits spherical symmetry to restore the lost dimension based on knowledge of the other two. If the original 3D density profile exhibits spherical symmetry in the yz -plane, the inverse Abel transformation yields [164]

$$n_\sigma(x, r) = -\frac{1}{\pi} \int_r^\infty dy' \frac{1}{\sqrt{y'^2 - r^2}} \frac{\partial n_{\text{col}}}{\partial y'}(x, y'), \quad (4.6)$$

with $r = \sqrt{y^2 + z^2}$.

In the experiment, clouds are not necessarily spherical symmetric but rather possess an elliptical-spherical symmetry. To apply the inverse Abel transformation for elliptical-spherical

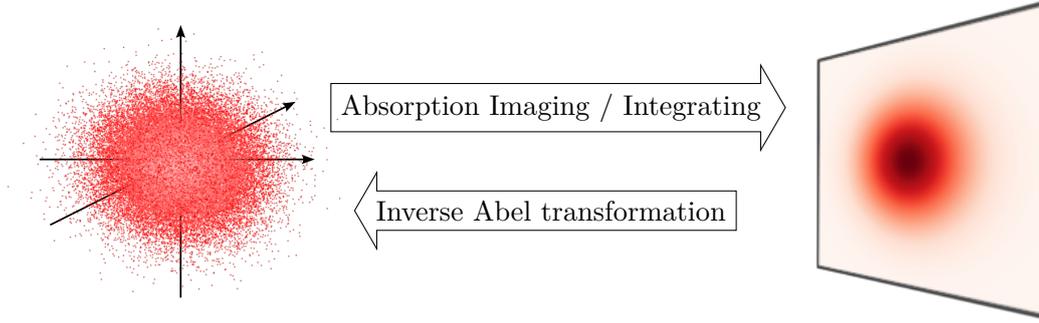


Figure 4.4: Principle of the inverse Abel transformation. In the experiment, a 3D cloud with density $n_\sigma(\mathbf{r})$ is prepared. By imaging the cloud, the 3D information gets lost as the density n_σ is integrated along the imaging direction z , yielding a 2D image of column densities $n_{\text{col}}(x, y)$. To reconstruct the original 3D density, an inverse Abel transformation is applied.

clouds, the radial coordinate is transformed to $r = \sqrt{y^2 + z^2 \cdot w_z^2/w_y^2}$, where w_z and w_y denote the cloud size in z and y direction, respectively. The aspect ratio w_z/w_y is measured via imaging from the side with Andor2 (refer to Section 4.1). With this transformation, the inverse Abel transformation adjusts slightly to [165]

$$n_\sigma(x, r) = -\frac{w_y}{w_z} \frac{1}{\pi} \int_r^\infty dy' \frac{1}{\sqrt{y'^2 - r^2}} \frac{\partial n_{\text{col}}}{\partial y'}(x, y'). \quad (4.7)$$

Note that σ is used as a subscript here, to clarify n_σ as the density per spin state σ . From here on, σ will generally be omitted, but n will still refer to the density per spin state unless stated otherwise.

Applying the inverse Abel-transforming

While the inverse Abel transformation is a straightforward technique for reconstructing the 3D density from a 2D projection, its application to real data requires careful considerations. First, the integral diverges at the initial value of $y' = r$. Second, the integral involves a numerical derivative of the measured column density n_{col} . Both of these aspects significantly amplify any noise in the experimental data. In order to mitigate this noise, the inverse Abel transformation is typically applied to 1D column density arrays, which average the column density along elliptical bins of equal column density. This algorithm has been developed by Andreas Kell [134] for our experiment.

Extracting the central density

With the algorithm applied, noise in the reconstructed 3D density is drastically reduced. However, since averaging the column density n_{col} along elliptical bins does not work at the trap centre, considerable noise remains in this region. Yet, the density at the trap centre is particularly important, as it is required to calculate the (homogeneous) Fermi energy (refer

to Section 2.4). To reduce noise at the centre, the central density is not directly deduced from the inverse Abel transformation, but is instead extrapolated from the density profile.

The extrapolation is performed using two distinct but equivalent approaches, illustrated in Fig. 4.5. The first approach interpolates the spatial density distribution with a centred Gaussian function, approximating the shape of the trapping potential. In the second approach, the position values are mapped to the potential of the calibrated trap (refer to Section 4.2.2). A linear interpolation between density and potential is then performed to obtain the centre density. This approach is justified by a good approximation of the centre region with a Thomas-Fermi distribution [8, 108, 134]. Both approaches have been used in this thesis. The latter approach has been used in Chapter 5, while the former one is applied to data in Chapter 7.

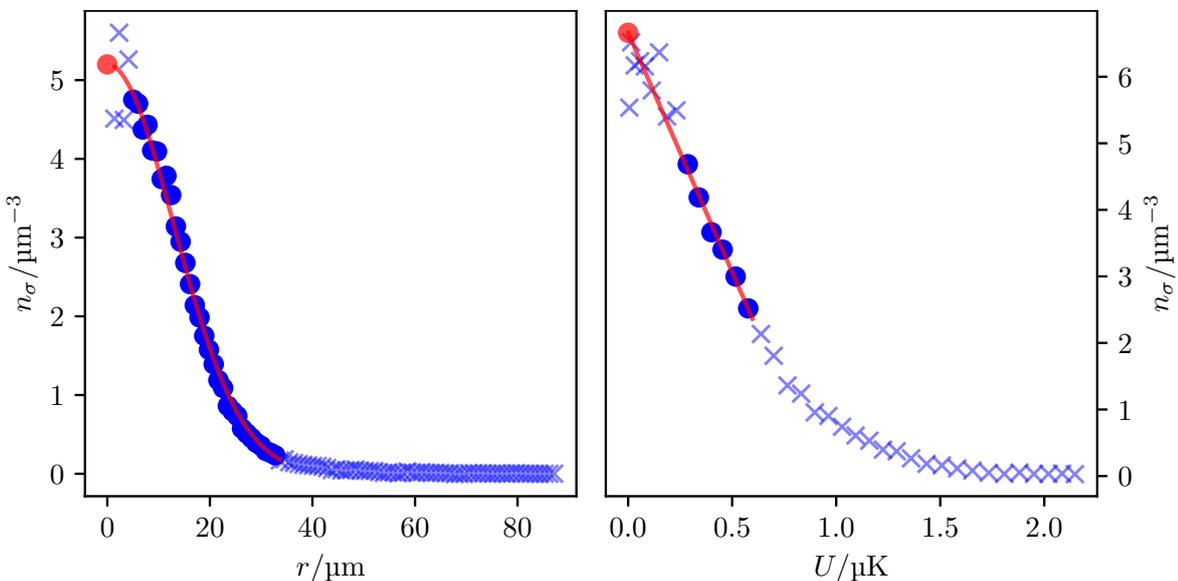


Figure 4.5: Interpolation of the centre density to mitigate noise. The figure shows two different interpolations (red lines) to determine the centre density (red dots). To this end, some values are omitted (crosses), while remaining values (blue dots) are used for the interpolation. **Left:** Spatial r density distribution n interpolated by a centred Gaussian function. **Right:** Different density distribution n interpolated by a linear function with respect to the trapping potential U .

Uncertainties due to anharmonicities of the trap

In the PhD thesis of Moritz Breyer [135], an extensive discussion is provided on the systematic uncertainty introduced by anharmonicities in the trap. The trapping confinement is created by two intersecting Gaussian dipole beams. In such a geometry, the equipotential lines stray away from ellipses in the outer regions of the trap, which disrupts the symmetry required for the inverse Abel transformation.

The uncertainty introduced by these anharmonicities has been estimated by simulating an unitary Fermi gas with its known EoS (refer to Section 4.2.2). By processing the simulated gas with the same procedure as introduced in this section, a small uncertainty of the central

density has been estimated to be around 5%. Since this also involves the reconstructed trapping potential in Section 4.2.2, the fitted temperature T is affected as well. Ultimately, a total error of approximately 5% has been concluded for the reduced temperature T/T_F .

4.2.2 Reconstruction of the potential

With the 3D density distribution in hand, the last remaining ingredient in determining the temperature is the reconstruction of a faithful trapping potential. As introduced in Section 3.2.6, the main trapping potential is created by a crossed optical dipole trap. Additionally, the curvature of the magnetic field from the Feshbach coils creates a confining potential in the horizontal plane, while the vertical direction is anti-trapped. Last but not least, gravity mgz adds to the total potential, which can be changed by applying additional magnetic gradients μB . Overall, the potential $U(\mathbf{r})$ is not given by a simple harmonic approximation but instead follows the combined expression

$$U(\mathbf{r}) = U_{\text{dipole, horizontal}}(\mathbf{r}) + U_{\text{dipole, vertical}}(\mathbf{r}) + U_{\text{Feshbach}}(\mathbf{r}) + mgz + \mu B. \quad (4.8)$$

The dipole trap provides the largest contribution to the trapping potential. In order to reconstruct the full trapping potential $U(\mathbf{r})$, several methods have been thoroughly explored and discussed in [134]. Shortly summarised, these methods rely on mapping the potential by probing the full extent of the dipole beam. In a first approach, the dipole beam has been directly imaged on the cameras used for absorption imaging. However, in this approach the beam passes through optics not designed for its wavelength, distorting the results. In a second approach, the power and waists of the dipole beams have been measured and interpolated to the atoms' position. From this, a relatively accurate potential could be retrieved. In a third approach, dipole modes (centre-of-mass oscillations) with increasing amplitude have been excited in the dipole trap in all three directions. Due to anharmonicities of the Gaussian beams, the trapping frequency is expected to decrease with increasing amplitude. However, diminishing signals at larger amplitudes and excessively large cloud sizes have prevented this method from reliably mapping the required potential range.

Ultimately, we have opted for mapping the potential using the known equation of state (EoS) at the critical point for a Fermi gas at unitarity [165]. The equation of state at unitarity is shown in Fig. 4.6. In the left plot, the relation between the chemical potential μ and temperature T with respect to the Fermi energy E_F and Fermi temperature T_F is displayed. The superfluid phase transition is measured to occur at $T_C/T_F = 0.167(13)$ and at a chemical potential of $\mu_C/\mu_F = 0.42(10)$, highlighted by a red cross [165]. Additionally, the reference also provides a measurement of the density n with respect to the ideal Fermi gas n_0 (refer to Eq. (4.4)) as a function of the logarithm of the fugacity $\log(z) = \mu/(k_B T)$. The red cross marks the critical point at $\mu_C/(k_B T_C) = 2.49$ and a critical density of $n_C/n_0 = 3.07$ [165]. Additionally, the panel shows the third and fourth order virial expansion of the equation of state at unitarity, which will be discussed further in Section 4.2.3.

To reconstruct the trapping potential, knowledge of the equation of state (EoS) is essential. Afterwards, inversion of the right panel of Fig. 4.6 yields the trapping potential in an LDA framework, with $\mu(\mathbf{r}) = \mu_0 - U(\mathbf{r})$. To gauge the EoS, the central fugacity $\mu_0/(k_B T)$ must be known. A convenient gauge point is the onset of superfluidity, where the condensate fraction

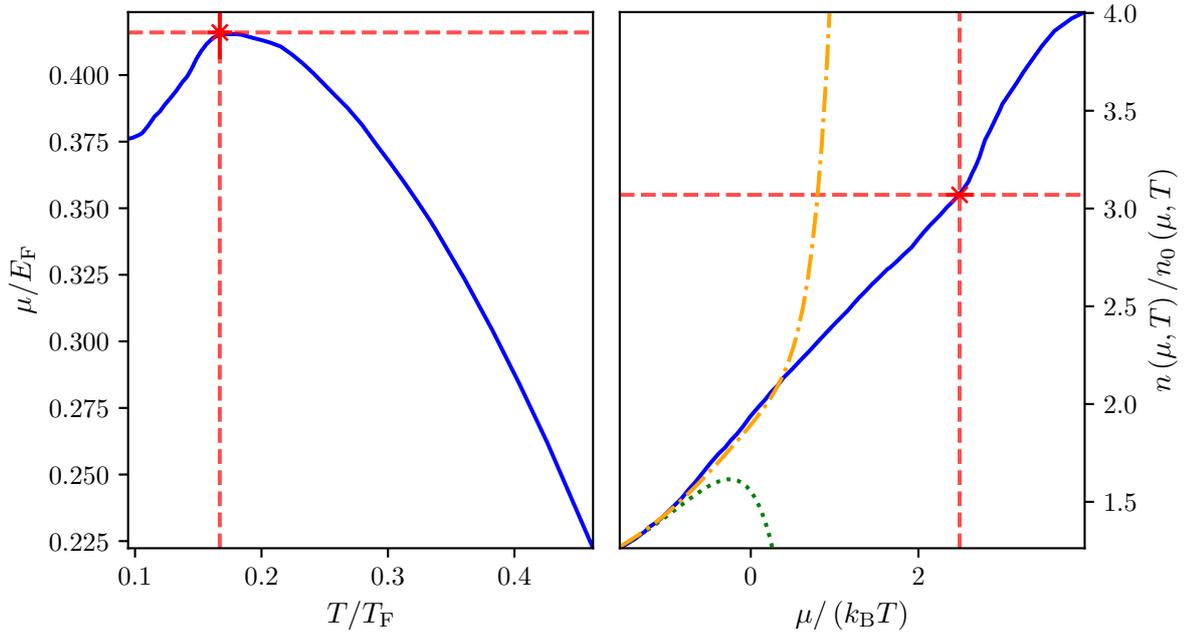


Figure 4.6: Equation of state of a unitary Fermi gas. The left panel shows the relationship between the chemical potential μ and temperature T/T_F . A phase transition (red cross) occurs at the maximum position of $\mu_C/E_F = 0.42(10)$, which takes place at a temperature of $T_C/T_F = 0.167(13)$. In the right panel, the density equation of state is shown, linking the fugacity $z = \exp(\mu/(k_B T))$ to the density of the unitary Fermi gas n in terms of the ideal Fermi gas density n_0 at the same (μ, T) . The red cross again marks the phase transition, while the green dotted and orange dash-dotted lines show the virial expansions of 3rd and 4th order, respectively. Data for the solid blue lines and critical points are taken from [165]. Virial coefficients are from [171].

turns zero $CF = 0$. At this point, the gas at the centre of the trap is exactly at the critical point $\mu_C/(k_B T_C) = 2.49$. In Section 4.3 and Section 4.4, techniques are introduced to detect the point where $CF = 0$. Finally, the trapping potential is deduced by measuring the 3D density profile of the superfluid critical gas at unitarity, followed by a subsequent inversion of the measured equation of state. This technique has been developed during the works on this thesis and is discussed extensively in [134].

In this thesis, two different traps are used for the conducted experiments. Both traps have been reconstructed using the described technique and are shown in Fig. 4.7.

4.2.3 Virial expansion of the equation of state - Temperature

During this section, we have already discussed the processing from column density n_{col} images to reconstructing the 3D density n_{σ} and the trapping potential $U(\mathbf{r})$. Now that the trapping potential $U(\mathbf{r})$ and the 3D density n are known, the temperature T can finally be determined by fitting the virial expansion of the equation of state (EoS) to the density distribution.

In Eq. (4.5), the Boltzmann distribution has been consulted to fit the density distribution of an ideal Fermi gas. However, in the BEC-BCS crossover studied here, the EoS is strongly influenced by interactions. In this regime, a virial expansion of the EoS yields a more accurate

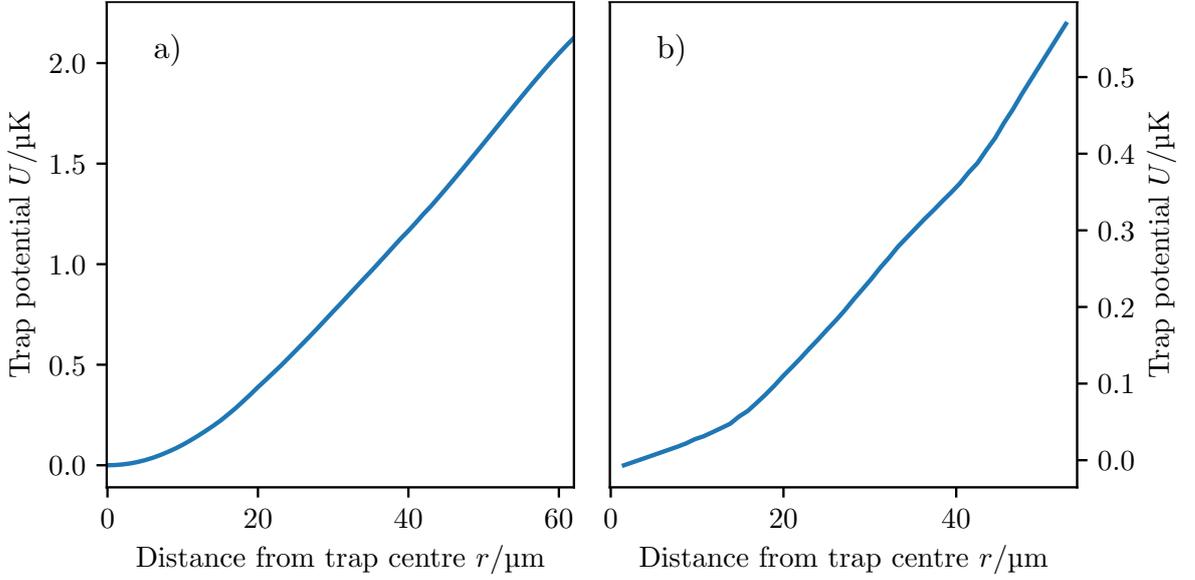


Figure 4.7: Reconstructed trap potentials. The left panel a) shows the trap potential for the machine learning project in Chapter 5, while the right panel b) shows the shallower trapping potential used in Chapter 7.

result.

The virial expansion separates the grand potential Ω into a power series of the fugacity $z = \exp(\mu/(k_{\text{B}}T))$, which reads [170, 172]

$$\Omega = -k_{\text{B}}T\tilde{V}\lambda_{\text{dB}}^{-3}\left(z + b_2z^2 + b_3z^3 + \dots\right), \quad (4.9)$$

with b_n being the virial coefficients, λ_{dB} the de-Broglie wavelength, T the temperature and \tilde{V} the volume of the homogeneous system. The density can be deduced from the power series as [170]

$$n_{\sigma} = -\frac{1}{\tilde{V}}\frac{\partial\Omega}{\partial\mu} = \lambda_{\text{dB}}^{-3}\left(z + 2b_2z^2 + 3b_3z^3 + \dots\right). \quad (4.10)$$

By comparison to the Boltzmann limit in Eq. (4.5), the first-order virial expansion with $b_1 = 1$ identifies as the Boltzmann equation of state.

The next virial coefficients b_n for the ideal Fermi gas can easily be calculated via [171]

$$b_n^{(0)} = (-1)^{n+1}n^{-\frac{5}{2}}. \quad (4.11)$$

However, these virial coefficients become much more complicated when interactions are introduced. To correct for interactions, the virial coefficients $b_n = b_n^{(0)} + \Delta b_n$ are given as corrections Δb_n to the coefficients of the ideal Fermi gas $b_n^{(0)}$. The corrections to the coefficients themselves now have a temperature dependence. In second order, the correction

is well-known and reads [173, 174]

$$\Delta b_2 = \begin{cases} \frac{e^{x^2}}{\sqrt{2}} [1 - \operatorname{erf}(|x|)], & \text{if } x < 0. \\ \sqrt{2}e^{x^2} - \frac{e^{x^2}}{\sqrt{2}} [1 - \operatorname{erf}(|x|)], & \text{otherwise,} \end{cases} \quad (4.12)$$

with $x = \lambda_{\text{dB}} / (\sqrt{2\pi}a)$.

The calculation of higher-order ($n \geq 3$) corrections Δb_n is much more involved, and various approaches have been studied to determine these corrections. Today, virial coefficients are known up to 3rd order on the BEC side and up to 5th order at unitarity and towards the BCS side [171, 172, 174, 175]. A comparison of the results has been discussed in detail in [134]. In the right panel of Fig. 4.6, the convergence of the 3rd and 4th order virial expansions to the measured density profile of the unitary Fermi gas is shown. While the expansion provides a good approximation in the low fugacity region $z < 1$ of $\mu/(k_{\text{B}}T) < 0$, the expansion diverges quickly in the high fugacity $z > 1$ region of $\mu/(k_{\text{B}}T) > 0$. Therefore, fitting the virial expansion to real data must focus on the low-density (low-fugacity) region, where $\mu/(k_{\text{B}}T) \ll 0$.

In order to determine the temperature of an in-situ image, the density profile is fitted to the virial expansion in Eq. (4.10) up to order m . Inclusion of up to m orders is sufficient if the thermometric values from the fit do not significantly change under the addition of the $m + 1$ order. An extensive study of the included orders and the resulting thermometric parameters is provided in [134]. In the determination of temperatures in Chapter 5, this involves a virial expansion of up to 3rd order, while samples furthest on the BCS side require the addition of the 4th order. The analysis of the new cooling technique in Chapter 7 operates solely on the BEC side and involves an expansion up to 3rd order.

In Fig. 4.8, an exemplary fit of the 3rd order virial expansion to the density profile of a sample on the BEC side is shown. As previously discussed, the virial expansion only describes the data in the low fugacity region of $\mu/(k_{\text{B}}T) \ll 0$, and thus we trim the fitting range accordingly. While this trimming is mainly important for unitarity and the BCS side, where $\mu_0 > 0$, the trimming also removes noise at the trap centre (refer to Section 4.2.1). Thus, the fitting region is confined to low densities. The exemplary fit yields a temperature of $T = 0.175(6) \mu\text{K}$ and a chemical potential of $\mu_0/k_{\text{B}} = -4.931(6) \mu\text{K}$.

4.3 Detecting Superfluidity via the rapid ramp technique

A key property in the BEC-BCS crossover is the condensate fraction. As discussed in Section 2.3, the onset of condensation coincides with the onset of the superfluid phase. In order to detect the superfluid phase transition or to probe excitations of the superfluid phase, the detection of the condensate fraction is an essential tool. The condensate fraction measures the amount of atoms participating in the macroscopic, condensed wavefunction N_{c} relative to the total number of atoms $N = N_{\text{c}} + N_{\text{t}}$ in the sample, where N_{t} represents atoms in the normal or thermal phase.

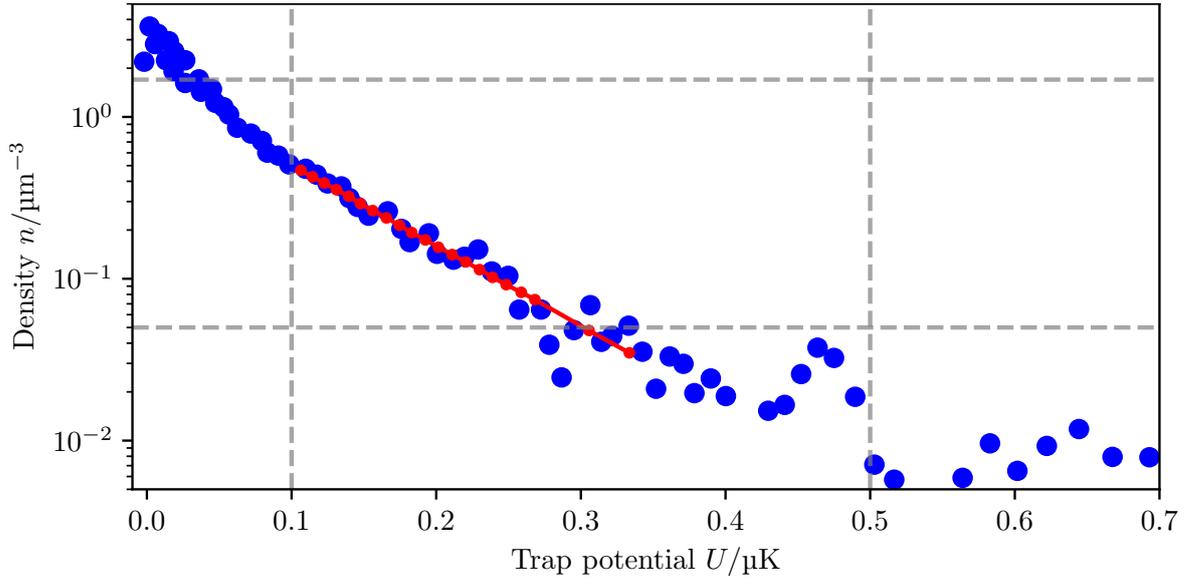


Figure 4.8: Fit of the virial expansion to the density. The plot shows a fit (red connected dots) of the virial expansion (3rd order) to the 3D density distribution (blue points) of the sample. Grey dashed lines highlight the trimmed fitting region. This example shows data used in Chapter 7 at an interaction parameter of $1/(k_F a) \approx 2.0$. The fit results in $T = 0.175(6) \mu\text{K}$ and $\mu_0/k_B = -4.931(6) \mu\text{K}$.

Condensation in the BEC limit

In the BEC limit of the crossover, the condensate fraction can easily be obtained through absorption imaging of an expanding cloud, a technique known as time-of-flight (“TOF”). As described in Section 4.1, this is achieved by allowing the cloud to expand for a duration of $T/4$ into a residual potential with frequency $\omega/(2\pi) = 1/T$. An absorption image taken after $T/4$ reveals the original momentum distribution [8]. Since the BEC occupies the ground state, the number N_c of composite dimers that participate in the BEC appear as a zero-momentum peak in the centre of the absorption image. All remaining dimers populate the available states according to Bose statistics (see Eq. (2.5)), while unpaired atoms follow Fermi statistics (see Fig. 2.1). The sum of those remaining atoms (atoms + dimers) N_t contribute to a diffuse thermal background. This is the smoking-gun of BEC, which provides a direct signature of condensation as a bimodal distribution of a condensate peak on top of a thermal background in images of the optical density. The bimodal distribution $f(x, y)$ in absorption images can be modelled as

$$f(x, y) = A_t \exp\left(-\frac{(x-x_0)^2}{2\sigma_x^2} - \frac{(y-y_0)^2}{2\sigma_y^2}\right) + A_c \left(1 - \frac{(x-x_0)^2}{w_x^2} - \frac{(y-y_0)^2}{w_y^2}\right)^{3/2} + C, \quad (4.13)$$

where C is an offset parameter.

Here, the first term represents the thermal background, while the second term comprises the macroscopic wavefunction of condensed atoms in the ground state. This function is fitted to optical density OD_0 images, and then the number of atoms in the condensate N_c and the

thermal background N_t are extracted by integrating over the corresponding terms

$$\begin{aligned} N_c &= \int dx \int dy A_c \left(1 - \frac{(x-x_0)^2}{w_x^2} - \frac{(y-y_0)^2}{w_y^2} \right)^{3/2} \\ N_t &= \int dx \int dy A_t \exp \left(-\frac{(x-x_0)^2}{2\sigma_x^2} - \frac{(y-y_0)^2}{2\sigma_y^2} \right). \end{aligned} \quad (4.14)$$

Finally, the condensate fraction CF is calculated by comparing the number of atoms in the condensate to the total number of atoms

$$CF = \frac{N_c}{N_c + N_t}. \quad (4.15)$$

Note that in the experiment, we only detect the atoms per spin state σ , but this does not affect the calculated condensate fraction CF .

Condensation in the BEC-BCS crossover

At unitarity and on the BCS side of the crossover, identifying condensation becomes more challenging. In these regimes, the condensate is also formed by zero-momentum pairs, i.e. Cooper pairs, but the bound state is a pure many-body effect that only forms due to the existence of the Fermi sea. If such a sample is expanding, the Fermi sea collapses due to the diminishing density, causing the bound state to disappear. Therefore, no zero-momentum peak of condensed pairs can be observed here.

In order to still detect pairing, correlations between images of the present spin states could be examined [176]. However, this approach is unfeasible with our experimental setup, as only one spin state can be imaged per experimental run.

Nevertheless, pairing still affects the momentum distribution of the atoms, as shown by the solution of the BCS Hamiltonian in Fig. 2.6. This influence, however, constitutes a very weak signature, as temperature and the presence of an inhomogeneous trap lead to similar changes in the momentum distribution. Albeit this obstructs the direct analysis of TOF images with a suitable model function, the remaining signature still proves to be detectable through machine learning techniques. These will be introduced in Section 4.4 and employed in Chapter 5.

Principles of the rapid ramp

Here, we introduce an alternative approach for detecting the condensate fraction across the entire crossover: the rapid ramp (RR) technique. The technique has been developed for the first observation of condensation throughout the BEC-BCS crossover in [38, 49]. In principle, the technique aims to recreate the bimodal distribution from Eq. (4.13) for arbitrary initial interactions.

To this end, the original pairs are projected onto tightly bound dimers via a ramp of the magnetic field, as illustrated in Fig. 4.9. These dimers, being two-body bound states, remain stable even at low densities. By ramping the magnetic field to the zero crossing of the scattering length a (refer to Section 2.2.2), the pairs are projected to a non-interacting

state of tightly bound dimers. Since this effectively turns off interactions, the momentum distribution is unaltered during expansion.

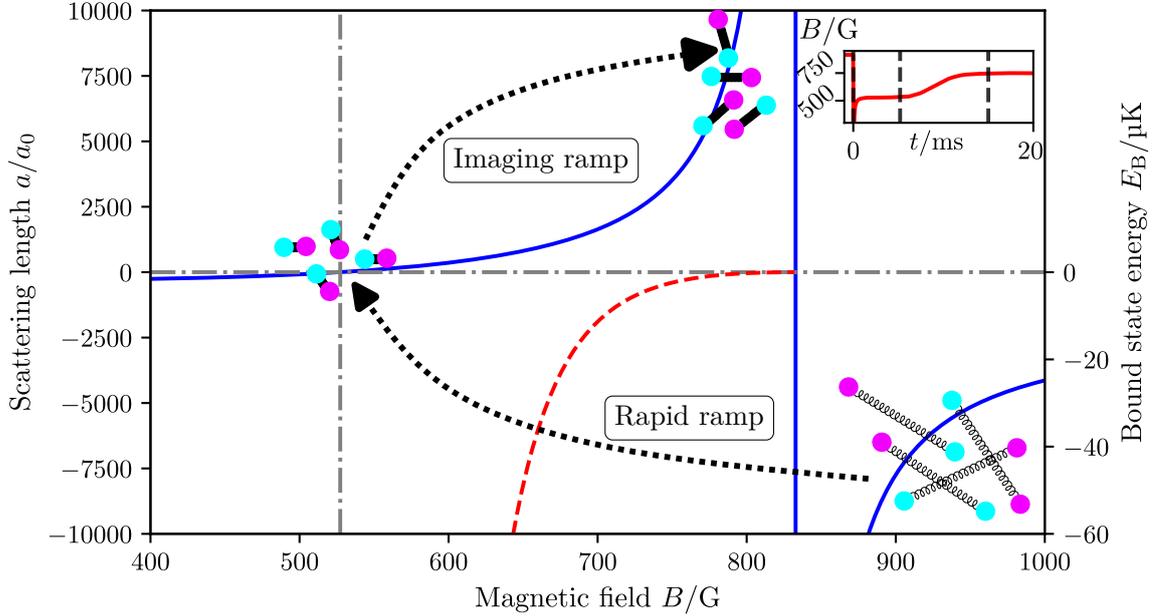


Figure 4.9: Principle of the rapid ramp. At the start of the free expansion, the magnetic field is ramped to the zero crossing of the scattering length (vertical dashed, dotted line). This “rapid ramp” projects pairs onto deeply bound dimers, preserving their momentum distribution. At the end of the expansion, the magnetic field is ramped back (“imaging ramp”) to a point where the bound state energy is small enough to simultaneously image free atoms and dimers. Data as shown in Fig. 2.3 for the $|12\rangle$ resonance [111]. The magnetic field ramps (measured via a current transducer) are shown in the inset in the upper right corner, with dashed lines marking the start of the expansion, the end of the free expansion at the zero crossing and the end of the whole expansion after 15 ms.

The ramp is initiated at the exact moment the trap is turned off and is performed fast enough to constitute a non-adiabatic change of the system, while still being slow enough to allow for the conversion of pairs to dimers. In the context of the Fermi gas, this constitutes a ramp faster than the relaxation time of the system, while still being slower than the much faster timescale of two-body physics [8]. At the end of the expansion time, when the density is already low and interactions play a less perturbative role, a second ramp returns the magnetic field close to unitarity. At this magnetic field, dimers and unpaired atoms can be imaged simultaneously, as the bound state energy is very small compared to the linewidth of the imaging transition.

Implementation of the rapid ramp

The rapid ramp (RR) has been integrated by the first two PhD students in the lab [40, 41]. During the works on this thesis, an improved circuit with better temporal resolution for initiating the ramp on a timescale of $1\ \mu\text{s}$ has been implemented by [135].

The complete ramp, comprising the rapid ramp to the zero crossing of the scattering length,

and the subsequent ramp to the imaging field, is performed over $T/4 = 15$ ms. Especially the rapid ramp is technically demanding, as the magnetic field must be changed within ~ 100 μ s from an initial magnetic field of 600 G – 1000 G to the zero crossing at ~ 527 G for the $|12\rangle$ mixture, or ~ 568 G for the $|13\rangle$ mixture. This rapid change is achieved by dissipating the magnetic field energy through an RC snubber circuit, while discharging the power supply via an additional IGBT. After a variable time τ_{off} , the dissipation through the snubber circuit is stopped, while the power supply discharge is stopped after a variable time $\tau_{\text{discharge}}$. In the meantime, the power supply set point is adjusted to match the magnetic field of the zero crossing in scattering length. The zero crossing is reached after ~ 260 μ s. After ~ 5.2 ms, the field is ramped back to the imaging (magnetic) field over a span of 10 ms. A full description of these steps and the experimental determination of the variable time constants τ_{off} and $\tau_{\text{discharge}}$ is found in [41, 135]. For faster experimental sequence setups, Andreas Kell has developed a deterministic model function to estimate good values for both τ_{off} and $\tau_{\text{discharge}}$ as functions of the initial magnetic field [134].

In the upper right corner of Fig. 4.9, the magnetic field ramp measured via a current transducer is shown. Note that the fast initial drop in the current transducer is much faster than the actual changing rate of the magnetic field at the atoms' position due to eddy currents in the chamber. The real magnetic field changing rate has been measured via the imaging detuning, which yields the stated duration of ~ 260 μ s to reach the zero crossing [40].

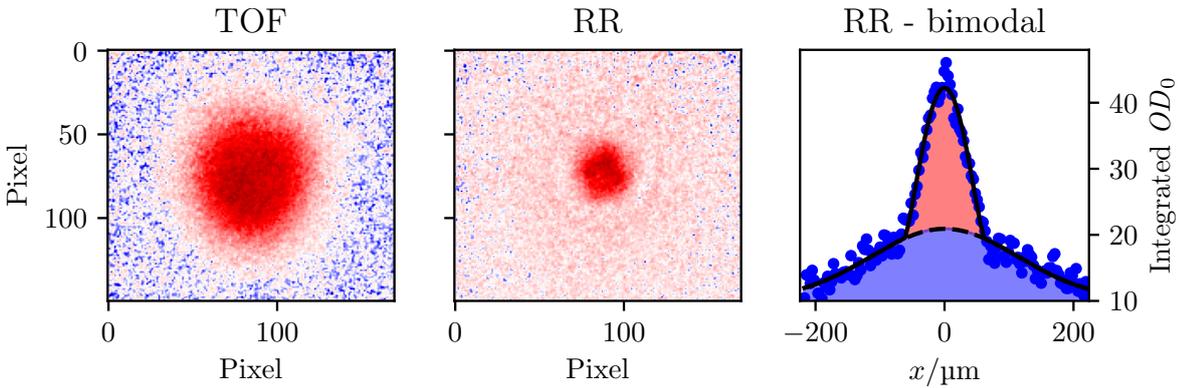


Figure 4.10: Absorption images (TOF) with and without the RR. Typical absorption images of a superfluid sample ($|12\rangle$ mix) after 5 ms TOF (left) and after 15 ms TOF with the RR technique (centre). Both images have the same experimental preparation and colour scale. After the RR, a sharp peak of the condensed zero-momentum pairs appears within a large thermal cloud surrounding it (centre). The bimodal distribution of the condensed peak (red shaded area) and thermal Gaussian background (blue shaded area) is more evident in the OD integrated along one direction (right), showing a condensate fraction of $CF \approx 0.26$.

After reaching the imaging field, an absorption image of the atoms reveals the condensate fraction CF . An exemplary absorption image in TOF with and without the RR is shown in Fig. 4.10. Note that the TOF image without the RR uses an expansion time of 5 ms, while the TOF image with the RR is taken after 15 ms of expansion. The RR is fitted to the model in Eq. (4.13), which yields the condensate fraction.

Limitations of the rapid ramp

The significance of the rapid ramp directly depends on how true the measured condensate fraction is to the real condensate fraction before the ramp. Ideally, the rapid ramp should conserve the momentum distribution, and the number of condensed pairs should remain constant. This would be true for a purely projective ramp of the magnetic field. However, as discussed, the actual magnetic field ramp is influenced by the exponential decay of the magnetic field due to eddy currents. As a result, the ramp’s timescale shifts from being purely projective to a more adiabatic ramp.

Moreover, it has been shown that the conversion efficiency of the rapid ramp from pairs to dimers depends on the initial magnetic field. Thus, the rapid ramp may not precisely reflect the exact condensate fraction before the ramp, but rather a fraction of it [50–52]. However, measurements have shown that the condensate formation time is significantly longer than the magnetic field ramp, ruling out the formation of newly condensed pairs during the ramp [166]. Thus, the rapid ramp remains the state-of-the-art technique for detecting condensation and changes in condensation across the crossover.

However, a precise determination of the onset of condensation may be obscured by perturbations introduced by the rapid ramp, such as by the detection of preformed pairs in a “pseudogap” state above the phase transition [51]. For this reason, new detection methods will be introduced in the following section.

4.4 Machine Learning

In this thesis, we employ neural networks as an alternative approach to detect the onset of superfluidity from TOF images without relying on the RR. Neural networks have already been applied to detect phase transitions in various system using both theoretical data [61–64] and experimental data [65, 66]. The initial application of neural networks to detect the onset of superfluidity in our experiment has been started by Martin Link [133]. During the works on this thesis, more effort has been made in combining the onset of superfluidity with a precise thermometry of the sample (refer to Section 4.2) [67, 134, 135]. A main result of this thesis constitutes the employment of an autoencoder neural network to detect the onset of superfluidity without using the RR at all [68]. This work, which has been started by [133], has been brought to completion in this thesis.

This section introduces the principles of neural networks and the learning algorithms they employ. Neural networks have been employed in several disciplines, including physics [177]. The networks used in this thesis are deep convolutional neural networks, which have been demonstrated to be well-suited for the task of image recognition [178]. In our case, the input images are TOF images, as shown in the left panel of Fig. 4.10. To detect the onset of superfluidity from these images, we use two distinct network architectures. The first one has several fully connected layers behind the convolutional layers to extract physical properties from the image, such as the condensation fraction CF . This network is trained in a supervised manner with labelled data, where the labels are retrieved from RR images providing the condensate fraction CF . The second architecture describes an autoencoder neural network with symmetric input and output layers. This network is trained to reproduce input data as accurately as possible, while a low-dimensional central layer forces the network

to learn a compressed low-dimensional representation of the input data [179]. After successful training, this low-dimensional representation can be extracted and analysed. Here, a second-order phase transition is found to be encoded as a change of trend in the low-dimensional distribution of TOF data.

In order to understand the working principle of these networks, this section starts by introducing a single artificial neuron. The connection of multiple of these neurons make up the artificial neural network. When each neuron in layer i is connected to every neuron in the subsequent layer $i + 1$, the layer is referred to as a fully connected or dense layer. Special attention is given to convolutional layers, which make up the first few layers of the networks in this thesis.

More details about the concepts of machine learning can be found in [133, 180, 181]. An excellent visual explanation of neural networks is available in [182].

4.4.1 Artificial neural networks

Artificial neurons are the fundamental entities of more complex neural networks [183]. A single artificial neuron is shown in Fig. 4.11. It comprises a number n of inputs x_i and corresponding weights w_i .

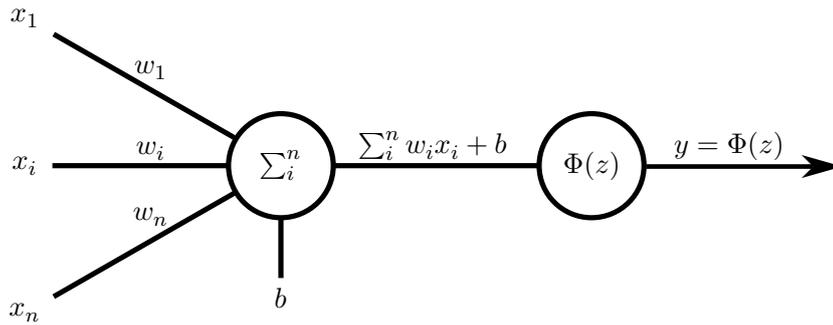


Figure 4.11: Concept of an artificial neuron. The sketch illustrates the workings of an artificial neuron with inputs x_1 to x_n and corresponding weights w_1 to w_n . An additional bias b is added to the weighted sum $z = \sum_i^n$. Finally, the output y is obtained by applying the activation function $\Phi(z)$ to the weighted sum.

The weighted sum of inputs is processed by a non-linear activation function Φ , yielding the output of the neuron [183]

$$y = \Phi\left(\sum_{i=0}^n w_i x_i\right) = \Phi(\mathbf{w}\mathbf{x} + b), \quad (4.16)$$

which incorporates a bias $w_0 x_0 = b$. In the final step, the weighted sum has been interpreted as the scalar product of the input vector $\mathbf{x} = (x_1, x_2, \dots)^\top$ and the weight vector $\mathbf{w} = (w_1, w_2, \dots)$. This will prove useful, as the functions of different layer architectures can be understood through linear algebra.

Hereafter, the artificial neuron will be simplified by a single quantity y defining its output value, connected to many inputs x_i , which are outputs y_i from the previous layer, with

weights w_i . For simplicity, the term “artificial” will be dropped, henceforth.

Activation function

The activation function Φ can be any differentiable function. In order to enable the neural network to approximate universal functions, the activation function is chosen to be non-linear [184, 185]. Popular choices for the activation function are listed in Fig. 4.12. The activation function is applied to the weighted input of $z = \sum_{i=0}^n w_i x_i$, and is usually centred around $z = 0$. Here, the role of the bias b becomes apparent as a shift $\sum_{i=1}^n w_i x_i > -b$ to the “activation threshold” of the neuron.

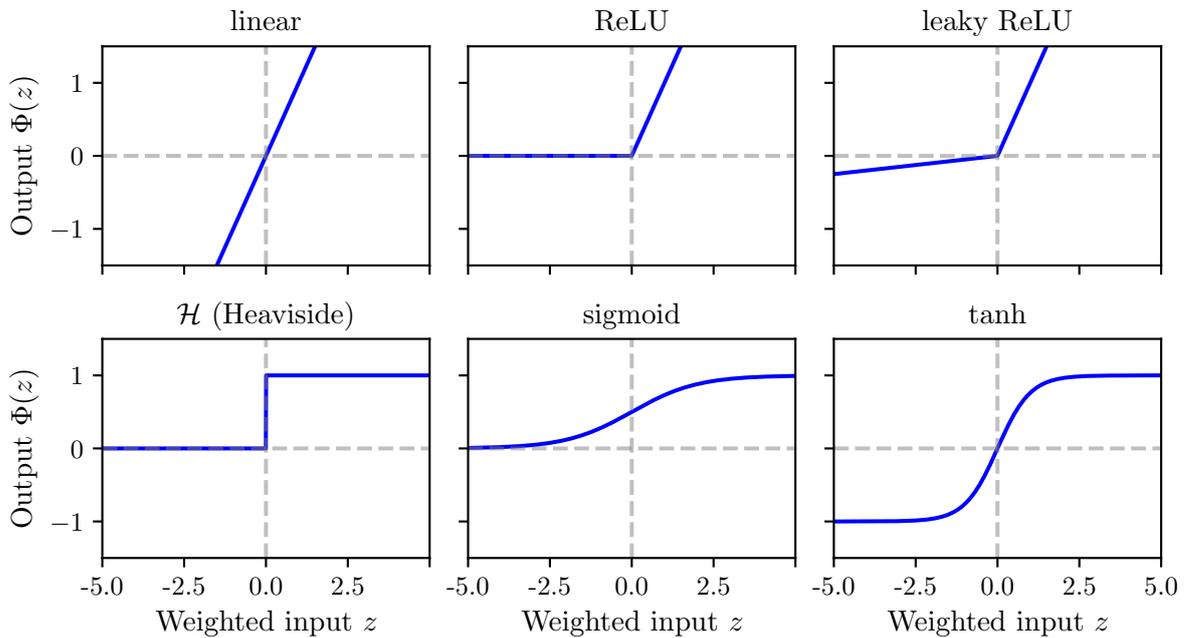


Figure 4.12: Overview of activation functions. The output of a neuron y applies an activation function $y = \Phi(z)$ to the weighted input z to this neuron (refer to Fig. 4.11). The figure presents various popular options for activation functions.

The most simple artificial neuron is created by employing the Heaviside \mathcal{H} function as its activation function. This constitutes a neuron that “fires” $y = 1$ if the input exceeds the threshold $z > 0$ or keeps “quite” $y = 0$ otherwise. Such an artificial neuron is known as the perceptron [186]. However, the Heaviside function has a major drawback: its gradient is zero everywhere. This prevents neural networks with the Heaviside function as the activation function from adjusting their weights via gradient propagation (refer to Section 4.4.2). Basically, such network are unable to be effectively trained.

Other popular choices for the activation function include the sigmoid function $f(z) = 1/(1 + e^{-z})$ [185] and the hyperbolic tangent function $\tanh(z)$. While \tanh has a greater rate of change, which can be beneficial for gradient descent, the output of the sigmoid function is limited to positive values, which may be useful for binary (yes or no) classification tasks. The rectified linear unit (“ReLU”) function $f(z) = \max(x, 0)$ has been shown to enable faster

convergence during training via gradient descent [187]. This comes naturally, as its gradient does not decay for large values of the input $z \gg 1$. However, its gradient vanishes if the weighted input turns negative $z < 0$. A modified version, the leaky ReLU, addresses this issue by introducing a small remaining positive gradient $\frac{\partial f(z)}{\partial z} = \epsilon$ for $z < 0$.

While the sole usage of purely linear activation functions in every layer prevents the neural network from becoming a universal function approximator [184], it can be necessary to employ linear activation functions in certain layers if the output range is desired to be unrestricted and linear. This is the case in the bottleneck layer of the autoencoder neural network in Section 5.5.

Deep neural networks

A neural network is created by connecting multiple artificial neurons across n distinct layers, as illustrated in Fig. 4.13. Each layer i comprises a number of neurons, highlighted by circles. Here, each circle represents a single neuron, as introduced in Fig. 4.11. The layers are stacked sequentially, starting with the input layer $i = 1$ and ending with the output layer $i = n$. If the network is made up of more than just the input and output layer, it is termed “deep”, with $n - 2$ hidden layers between the input and output layer [181]. The dimensionality of the input layer defines the dimensionality of the input data, such as the number of pixels in an image. For instance, for an image of $100 \text{ px} \times 100 \text{ px}$ pixels, this already constitutes an input dimensionality of 10000. The dimensionality of the output must match the number of predicted quantities, such as the number of classes in a classification task or the number of fitted values.

Each layer $i \neq 1$ has connections of its neurons to the neurons of the previous layer $i - 1$. These connections are quantified by weights w , which depends on the architecture of the layer. Each architecture serves a distinct purpose and use case. Stacking many of these layers constitutes the final neural network.

Dense layers

A dense layer describes a fully connected layer i , where each neuron in layer i has a connection to each neuron in the previous layer $i - 1$. Thus, each neuron j in layer i receives a weighted sum of every neuron x_k^{i-1} from layer $i - 1$, which reads $z_j^i = \sum_k w_{jk}^i x_k^{i-1}$. The weights w_{jk}^i of layer i can be understood as matrix elements of the weight matrix \mathbf{W}^i . Therefore, the operation of the dense layer can be formulated as a matrix multiplication of \mathbf{W}^i with the (input) neuron vector $\mathbf{x}^{i-1} = (x_1^{i-1}, x_2^{i-1}, \dots)^\top$, yielding the output vector [181]

$$\mathbf{y}^i = \Phi(z^i) = \Phi(\mathbf{W}^i \cdot \mathbf{x}^{i-1} + \mathbf{b}^i), \quad (4.17)$$

with \mathbf{b}^i being the bias vector of layer i . The activation function Φ is applied to every element z^i .

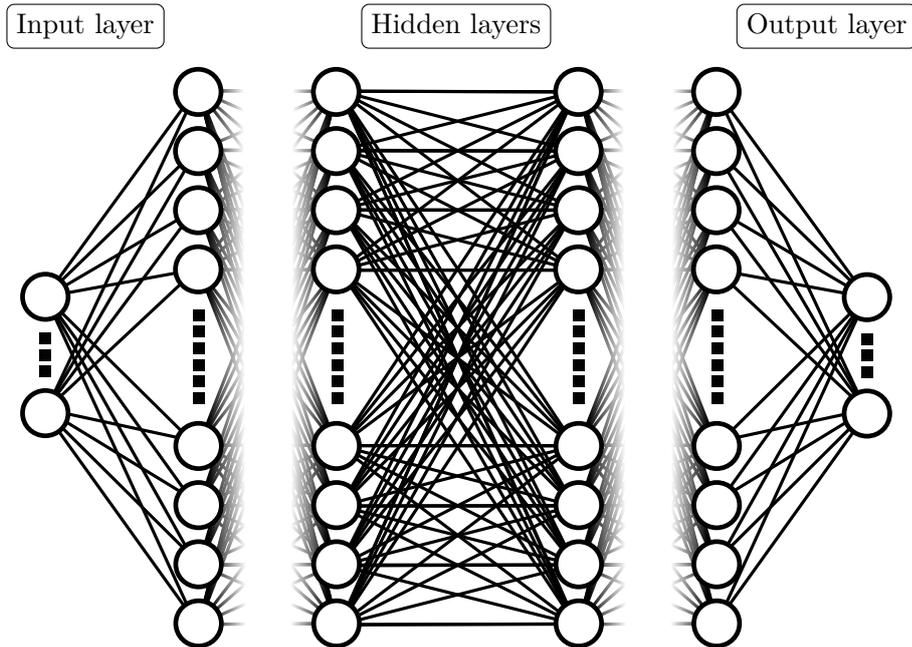


Figure 4.13: Sketch of a deep neural network. The sketch shows an arbitrary deep neural network. An n -dimensional input is processed by k consecutive hidden layers, yielding an m -dimensional output. Each circle represents a neuron, while only the input and output neurons are accessible. A neuron in layer i can be connected to every neuron in layer $i + 1$ (“dense layer”).

Convolutional layers

Convolutional layers describe an alternative layer architecture to dense layers. In a dense layer, every neuron of layer i is connected to every neuron in layer $i - 1$ with a distinct weight w_{jk}^i . A convolutional layer i uses a set of n kernels or masks, which are moved across the input layer $i - 1$ to create a weighted input z for every kernel entity and position. Each kernel m has the same size of $x \times y$ weights w_{mjk}^i . Thus, a convolutional layer has only $n \times x \times y$ different weights, shared across the entire input layer $i - 1$ to compute the weighted sums. Since the convolutional layer has n kernels, the output of the layer acquires a third dimension.

The operation of a convolutional layer on an input layer is illustrated in Fig. 4.14. In this example, the convolutional layer has $n = 3$ kernels, each of size 3×3 . The kernels are moved across the input layer in steps defined by the hyperparameter “strides”. For every position (α, β) of kernel m , an output value y_m is computed by applying the activation function to the weighted sum [181]

$$y_m = \Phi \left(\sum_{mjk} w_{mjk} x_{\alpha+j, \beta+k} \right). \quad (4.18)$$

Padding along the edges of the input layer is added to increase the size of the output layer, preserving the size of the input if desired. Neurons in the padding area have a value of zero. In the example, the output of the convolutional layer has gained a third dimension due to the set of $n = 3$ kernels.

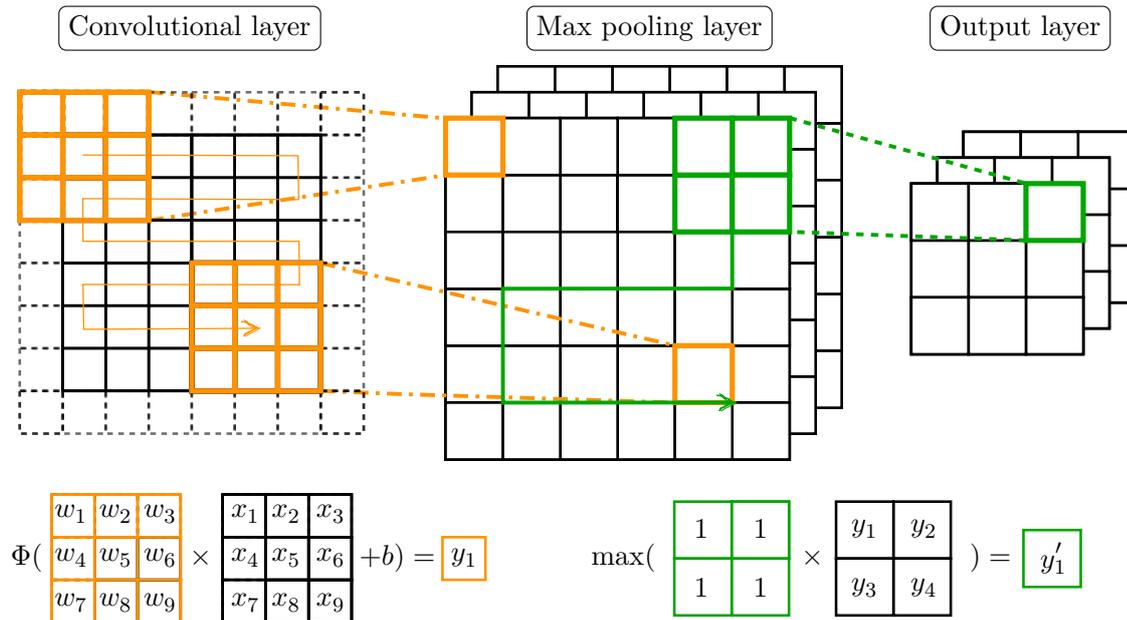


Figure 4.14: Convolutional plus max pooling layers. The sketch illustrates the operation of a 2D convolutional layer followed by a 2D max pooling layer. First, the convolutional layer is applied to the input data (solid 6×6 grid). To this end, a kernel (orange 3×3) with weights w_i moves (orange arrow) across the input layer according to the “strides” parameter. Additional padding (dashed outer squares) filled with zero values can be added to get the same input and output layer shape. The dimension of the output layer can be increased with the number of kernels (3 in this sketch). After the convolutional layer, a max pooling layer is employed, reducing the output size by fetching the maximum value along another mask (2×2 , green squares).

Convolutional layers are usually followed by a downsampling layer, implemented here as a max pooling layer. Max pooling layers work by moving a mask of size $x \times y$ across their input layer, with a step size equal to their mask size in each dimension. At each position, the maximum value within the mask is transferred to the output, effectively downsampling the input layer. This process is also sketched in Fig. 4.14.

Convolutional neural networks (CNNs) are inspired by the discovery that in animals, such as cats and monkeys, certain neurons are activated by specific regions in their receptive fields [188]. This concept is mirrored by the architecture of a convolutional layer, where kernels scan across the input layer [189]. By including subsequent max pooling layers, CNNs become highly effective at highlighting features in images. Thanks to weight sharing in convolutional layers and the downsampling in max pooling layers, CNNs require significantly fewer parameters, enabling faster computation and reduced memory requirements.

4.4.2 Supervised learning

Until now, we have just introduced potential architectures of neural networks. These networks are made up of a manifold of neurons (Fig. 4.11), arranged in layers and connected through

distinct weights. It is easy to understand that a unique stack of a few layers, with properly chosen weights, can reconstruct any function [189]. However, as we have seen during the introduction of dense and convolutional layers, the amount of parameters of such a network quickly becomes substantial, often in the range of 10000 to 1000000, or possibly even more. For a neural network tasked with interpreting images of $100 \text{ px} \times 100 \text{ px}$, this would constitute a function with an input dimension of 10000 and a highly convolved hidden layer structure, possibly containing between 10000 and 1000000 parameters. Consequently, such a network becomes extremely sensitive to changes in either input values or network parameters, such as the weights.

If the weights are randomly set, the network will produce random outputs for any input. In order to map inputs to desired outputs accurately, the weights must be carefully tuned. This has posed a tremendous technical challenge until the development of a technique called “backpropagation” [190]. Backpropagation allows the neural network to iteratively adjust its weights according to the deviation between its output and the desired output. This enables deep neural networks to become universal function approximators [184, 185].

Loss function

The discrepancy between the network’s output and the desired output is quantified by a loss function. A popular choice for the loss function L between the network output y and the desired output \hat{Y} is the mean squared error

$$L = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{Y}_i|^2, \quad (4.19)$$

for a dataset of N entries. The desired outputs \hat{Y}_i are manually determined “labels” of corresponding input data \hat{X}_i . Thus, the dataset with N entries of input data \hat{X}_i and labels \hat{Y}_i is called the “training dataset”. Usually, the training dataset is split into batches of $m \leq N$ data points, to decrease memory requirements and to optimise the weight adjustments [183, 191]. In the context of gradient propagation, this is called “stochastic gradient descent” [181].

Gradient propagation

Gradient propagation, or backpropagation, describes the adjustment of parameters in the network with the aim of minimising the loss function. When the loss function is minimised, the network has effectively learned to map input data to outputs that closely resemble the labels. To approach this minimum, the weights w_{jk}^i and biases b_j^i of the network are updated towards a smaller value of the loss function, which is mathematically expressed in terms of gradients

$$\begin{aligned} w_{jk}^i &\rightarrow w_{jk}^i - \eta \frac{\partial L}{\partial w_{jk}^i} \\ b_j^i &\rightarrow b_j^i - \eta \frac{\partial L}{\partial b_j^i}, \end{aligned} \quad (4.20)$$

where η is a hyperparameter known as the “learning rate”. A small learning rate updates the parameters in small increments, approaching the minimum very slowly. While this increases the sampling resolution of minima in the loss function, effectively making it more precise, it increases the risk of sticking to local minima in the hyperspace of the loss function. Much higher learning rates mitigate this problem, but have trouble to converge, which may lead to instability. In Chapter 5, we employ the Adam optimiser [192] during training, which adaptively adjusts the learning rate by calculating a moving average of the gradients from previous iterations, balancing convergence speed and stability.

The gradients are not trivial, but can be calculated by applying the chain rule. For the output layer $i = n$, this yields

$$\begin{aligned}\frac{\partial L}{\partial w_{jk}^n} &= \frac{\partial L}{\partial y_j^n} \times \frac{\partial y_j^n}{\partial w_{jk}^n} = \frac{\partial L}{\partial y_j^n} \times \frac{\partial y_j^n}{\partial z_j^n} \times \frac{\partial z_j^n}{\partial w_{jk}^n} \\ \frac{\partial L}{\partial b_j^n} &= \frac{\partial L}{\partial y_j^n} \times \frac{\partial y_j^n}{\partial b_j^n} = \frac{\partial L}{\partial y_j^n} \times \frac{\partial y_j^n}{\partial z_j^n} \times \frac{\partial z_j^n}{\partial b_j^n}.\end{aligned}\tag{4.21}$$

Each term in these expressions can be calculated straightforwardly. The first term is the derivative of the loss function, with the loss function being the mean squared error. The second term represents the derivative of the activation function $y = \Phi(z)$. This is why activation functions in neural network are chosen to be differentiable, as discussed in Section 4.4.1. Some activation function in Fig. 4.12 exhibit a kink at $z = 0$; however, the gradient for such functions is defined to be zero at this point. The last term denotes the derivative of the weighted input z with respect to the weights, which is just the input value along the weight.

In deeper layers $i < n$, the calculation of the gradient is more complicated but can be handled through the successive application of the chain rule. The calculation is similar to Eq. (4.21), but the derivative of the loss function L with the output y_j^i is expanded by applying the chain rule again

$$\frac{\partial L}{\partial y_j^i} = \frac{\partial L}{\partial y_p^{i+1}} \times \frac{\partial y_p^{i+1}}{\partial y_j^i} = \frac{\partial L}{\partial y_p^{i+1}} \times \frac{\partial y_p^{i+1}}{\partial z_p^{i+1}} \times \frac{\partial z_p^{i+1}}{\partial y_j^i},\tag{4.22}$$

where the final derivative can be replaced by the weight w_{jp}^i between layer i and $i + 1$. Finally, the gradient of $\frac{\partial L}{\partial w_{jk}^i}$ can be calculated by further applications of the chain rule for the term $\frac{\partial L}{\partial y_p^{i+1}}$ until reaching the output layer. Given the step-by-step backpropagation of these calculations, the method is aptly named “backpropagation”.

4.4.3 Unsupervised neural networks

In supervised machine learning, networks are trained on manually labelled datasets with pairs of inputs $\hat{\mathbf{X}}$ and corresponding labels $\hat{\mathbf{Y}}$. Alternative training methods are summarised under the umbrella term of “unsupervised machine learning”. As the name suggests, these are methods of inferring structure and information from data without the creation of a labelled dataset.

Unsupervised learning has diverse applications, such as image denoising, clustering data

based on similarity, compressing the dimensionality of the data representation and generating new data from this representation via diffusion, to name a few examples [181, 193, 194]. A prominent example of a model using unsupervised learning is the image diffusion model *Stable Diffusion* by *Stability AI* [55, 195]. This model combines a text encoder with a variational autoencoder (“VAE”) and a subsequent denoising network to generate images from a low-dimensional representation.

Autoencoder

Autoencoders are a subclass of unsupervised learning techniques. The core idea of the autoencoder is to generate a low-dimensional representation of a dataset, such as images [179]. To this end, the autoencoder represents a deep neural network shown in Fig. 4.13, but with the same dimensionality of its input and output layer. Within the hidden layers, one specific layer – known as the “bottleneck layer” – has a significantly smaller dimensionality (neuron number). By training the network via gradient propagation to reconstruct the input images as close as possible (effectively approximating the unity function), the network is forced to learn a low-dimensional representation of the data structure due to the bottleneck [54, 181]. This low-dimensional representation in the bottleneck is called the “latent space”, representing an encoding of input data, which contains as much information about the input as possible, such that the second half of the autoencoder can decode and reconstruct images back from the latent space.

Autoencoders are very similar to an alternative technique, called principal component analysis (“PCA”) [196]. PCA extracts a low-dimensional space by linearly transforming the input space, such that the new axes are sorted after highest variance of data. However, unlike PCA’s linear transformations, autoencoder can perform complex, non-linear transformations, yielding richer, non-linear encodings of data patterns.

The encoded data can be used to associate images with distinct classes based on their distribution in the low-dimensional latent space, e.g., images of dogs might be clustered in one region, while cats populate a disconnected region. Small dogs like chihuahuas might creep into the region of cats, but the autoencoder should be able to infer a trend from small to large dogs, that are very distinct from cats. Thus, this layout allows the model to infer information about a smooth “phase” transition – like from small to large dogs – based on trends in the low-dimensional latent space.

A similarly featureless phase transition is given by the second-order phase transition that is the superfluid phase transition in the BEC-BCS crossover. In this thesis, we employ an autoencoder neural network on TOF images of samples throughout the crossover for temperature above and below the phase transition. After successful training of the network, the autoencoder enables us to observe the onset of superfluidity as a distinct change in trend within its low-dimensional latent space.

4.5 Superimposing a fast magnetic field

In the experiment, we introduce perturbations into the system using an additional, homebuilt small magnetic field coil. This coil has specifically been designed with the objective of changing the magnetic field faster than the fastest timescale of the system, which is the Fermi

time $\tau_F = \hbar/E_F$ [8]. The typical timescale given by the Fermi time in our experiment is on the order of $5 \mu\text{s} - 10 \mu\text{s}$. For this purpose, a small coil comprising two concentric coils has been designed and built by [134, 163], which is depicted on the right side of Fig. 4.15. This coil is often referred to as the “quench coil”.

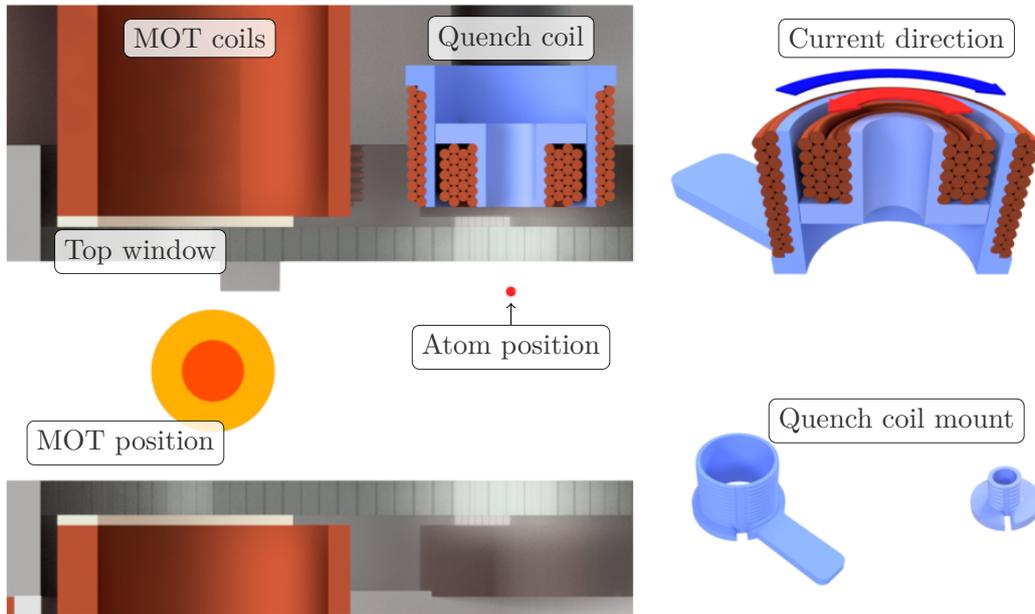


Figure 4.15: Setup of the quench coil. The render shows the relative position of the quench coil with respect to the atoms’ position (set by the dipole trap) and the MOT position (defined by the MOT coils). For scale, the atom position is just 3 mm below the top vacuum window. On the right side of the figure, the 3D-printed mount and winding orientation of the quench coil are shown. Figure adapted from [134].

This coil produces a finite magnetic field with a nearly zero gradient at the atoms’ position by using a design with a larger coil and a counter-winding smaller coil housed in a compact 3D-printed mount. The result is a low inductance and magnetic field range, enabling fast magnetic field changes. To retain a strong enough magnetic field at the atoms’ position, the coil is positioned as close as possible ($\sim 6.5 \text{ mm}$) to the atoms, directly on top of the main chamber’s upper viewport, as illustrated in Fig. 4.15. The different coil sizes ensure a small magnetic field while largely cancelling the gradient.

Calibration and alignment

As already discussed, the coil is designed to have a small remaining magnetic field at the atoms’ position while the gradient is mostly vanished. This is not true for the whole magnetic field, but only for a small volume at a specific distance from the coil.

To this end, the magnetic field has been calibrated using the same method as for the magnetic field of the Feshbach coils, discussed in Fig. 3.4. For different DC currents of the coil, the resonance position of the $|1\rangle \rightarrow |2\rangle$ transition has been tracked, revealing the

magnetic field of the coil. Via this method, a calibration factor of $0.874(8)$ G/A has been determined.

While the coil has been designed to have a zero gradient at the atoms' position if moved as close to the viewport as possible, it has been found that a small gradient along the (vertical) z axis of the coil remains. This gradient has been compensated for by adding counteracting gradients with otherwise unused coils.

The gradient in the yz -plane, perpendicular to the coil axis, has been eliminated by adjusting the coil position within this plane. If a gradient remains, atoms that are released from the dipole trap experience strong deflection due to the gradient. By aligning the coil properly, sloshing is minimised.

4.5.1 Quench of the magnetic field

The coil can operate in two distinct modes using separate driver circuits. In the first mode, a rapid quench of the magnetic field is achieved using the driver circuit detailed in [134, 163]. This is accomplished by swiftly switching off the current (within 30 ns) via a high-speed power MOSFET driver and a power MOSFET, followed by the rapid dissipation of remaining magnetic field energy through an RC snubber circuit. The dissipation rate determines the quench timescale, measured with a pick-up coil to be approximately $2.6 \mu\text{s}$. This quench speed is indeed faster than the Fermi time of $5 \mu\text{s} - 10 \mu\text{s}$, as required.

In this mode, magnetic field quenches of up to $35.0(3)$ G in amplitude are performed, corresponding to a maximum current of 40 A. The coil's magnetic field adds to the "final" magnetic field produced by the Feshbach coils. By turning the coil off, the magnetic field is quenched back to this "final" magnetic field. Since the quench is faster than the Fermi time, the quench yields a non-adiabatic change of the interaction parameter $1/(k_F a)$, calculated from the corresponding change in the scattering length a (refer to Fig. 2.3). This allows for the excitation of the Higgs mode, which will be explored in Chapter 6.

4.5.2 Modulation of the magnetic field

The second mode of operation employs an amplifier circuit [134] that enables the modulation of the magnetic field. To this end, an arbitrary function generator[†] provides a sinusoidal signal to the amplifier, which modulates the magnetic field accordingly. The amplifier is designed for an output current amplitude of up to 10 A at frequencies of up to 40 kHz. Higher frequencies are possible at a reduced output amplitude.

The calibration of the amplifier has been performed similarly to the DC mode explained in Section 4.5. Instead of calibrating the field via the known $|1\rangle \rightarrow |2\rangle$ transition frequency, the detuning to the imaging frequency of the D_2 transition (refer to Section 3.1.1) has been mapped to the magnetic field. Since imaging only takes a few microseconds, this method allows to resolve the magnetic field modulation up to frequencies on the order of ~ 100 kHz. While effective, this method requires substantial time for extensive data collection, and so it only serves as a reference at a few selected frequencies.

[†]GW Instek - AFG-2225

For broader parameter calibration, spanning input frequency and voltage of the function generator, we use the amplifier circuit's monitoring port[‡] to measure the coil current with a translation factor of 0.25 V/A. Considering the previously stated calibration factor of 0.874(8) G/A, this allows calibration of the output magnetic field amplitude G_p relative to the input amplitude V_{pp} (peak-peak).

In [134], a frequency range of up to 200 kHz has been analysed. The more intricate calibration via the imaging detuning probes the real field more carefully. This has revealed a frequency dependent correction factor to the monitoring-port calibration. The correction is considered for all calibrations, henceforth.

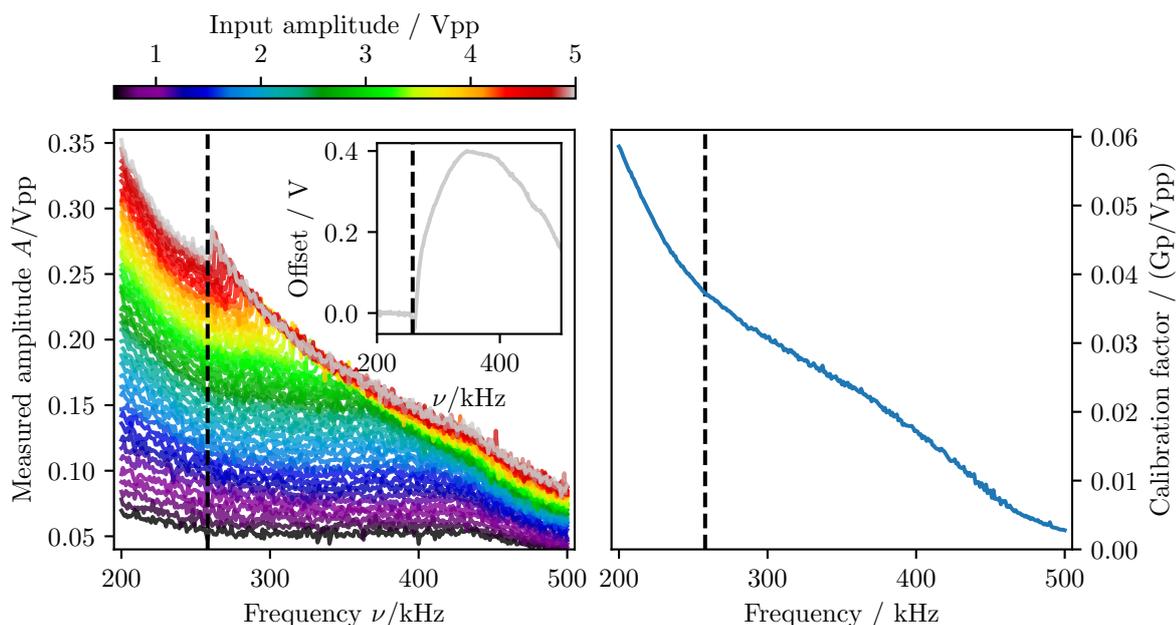


Figure 4.16: Amplitude calibration via the current monitor of the amplifier circuit for high frequencies. **Left:** Calibration data obtained from the current monitor A/V_{pp} (peak-peak) as a function of input amplitude and frequency of the function generator. **Right:** At each frequency, a linear fit of measured amplitude versus input amplitude provides the calibration factor G_p/V_{pp} . Note that this factor determines the magnetic field amplitude (zero-peak), and thus incorporates the magnetic field vs. current calibration and the frequency correction from [134]. The dashed line in both panels, as well as in the inset, is at 258 kHz. **Inset:** Above this frequency, the calibration becomes highly non-linear, which is mostly visible in the emergence of a DC offset field.

Here, we extend the calibration analysis to the > 200 kHz range, necessary for dissociating molecules in the far BEC regime in Chapter 7. Fig. 4.16 shows the calibration from 200 kHz to 500 kHz. This calibration smoothly connects to the existing calibration for frequencies < 200 kHz [163].

For frequencies larger 258 kHz, a notable offset or DC component appears in the output signal, dependent on the input amplitude. This offset has been confirmed as genuine, not an artifact of the monitoring output, by dissociating composite dimers at magnetic fields that

[‡]AD8429ARZ

correspond to bound state energies E_B/h beyond > 258 kHz (refer to Section 2.2.3). Here, it has been observed, that by increasing the input voltage to the amplifier, the dissociation threshold frequency shifts to lower values, which is caused by an (DC) increase in the magnetic field of similar strength as seen in the current monitor. The effect probably arises due to a sharp reduction in the common-mode rejection ratio of the amplifier circuit[§] for frequencies beyond 100 kHz. To prevent the offset from distorting the results in Chapter 7, no frequencies larger than 258 kHz have been studied.

[§]OPA549T

Machine Learning the Superfluid Phase Transition

This chapter presents the results of the two projects aimed at detecting the superfluid phase transition in a strongly interacting Fermi gas through advanced image recognition using machine learning techniques. To this end, two distinct neural networks are applied to TOF images, as introduced in Section 5.4 and Section 5.5.

The TOF images have been recorded in a $|12\rangle$ mixture at various magnetic fields, corresponding to interaction parameters $1/(k_F a)$. To detect the phase transition at each magnetic field, images have been taken at temperatures above and below the superfluid phase transition temperature T_C/T_F , achieved by a controlled heating of the sample. This approach yields a critical heating time τ_C for each magnetic field sequence, indicating when the sample transitions to the “normal” state. By repeating the experiment around the critical heating time, the corresponding temperature T_C and Fermi Energy T_F are determined via thermometry on in-situ images of the samples (refer to Section 4.2). Ultimately, this establishes the phase boundary between the superfluid and normal states within the strongly interacting regime $-1 < 1/(k_F a) < 1$.

The first method, detailed in Section 5.4, uses a deep convolutional neural network (refer to Section 4.4.1) trained to predict the condensate fraction directly from TOF images. For supervised learning, labels have been created by measuring the condensate fraction via the rapid ramp technique (refer to Section 4.3). By excluding rapid ramp labels near the phase transition, where the rapid ramp becomes unreliable, the network learns to predict the onset of condensation. After training, the critical heating time is identified as the point where the predicted condensate fraction vanishes. This project has originally been initiated by Martin Link [133]. During the works on this thesis, a more reliable thermometry has been developed [134, 135] (refer to Section 4.2). Additionally, the neural network model has been retrained following modifications to the experimental apparatus, which has been optimised through transfer learning [197–199]. This process also required the acquisition of new datasets for the thermometry and network predictions, i.e. the TOF inputs and RR labels. Now, this work is published in [67].

In the second approach, presented in Section 5.5, the superfluid phase transition is detected using unsupervised machine learning. Here, an autoencoder neural network (see Section 4.4.3)

is trained to reconstruct TOF images without additional inputs, such as RR labels or thermometric parameters. Due to the presence of a “bottleneck” layer, the network is forced to find a low-dimensional representation of input images that suffices to decode the data back into an output that closely resembles the original images. After successful training, the decoder part of the network is stripped off, allowing access to the low-dimensional representation of the input dataset in the “latent space”. In this latent space, the network learns to organise the data by temperature and the interaction parameter, solely based on the momentum distribution in TOF images. While the latent space does not explicitly separate images into the normal or superfluid phase, the second-order superfluid phase transition is indicated by a change in trend within the latent space distribution at specific heating times. From these trend-changing positions, a critical temperature can be inferred. By comparing the critical temperature with the established phase boundary from Section 5.4, the trend-changing position is identified as the superfluid phase transition. This project has been initiated by Martin Link [133] as well. In this thesis, the analysis has been expanded and employed to the new datasets from the supervised analysis, yielding a phase diagram solely from TOF images. The results of this project represent a primary goal of this thesis and have now been published [68].

Outline of this chapter

The sections of this chapter are outlined as follows:

- Section 5.1: In the first section we briefly discuss conventional methods and available signatures for detecting superfluidity.
- Section 5.2: The limitations of conventional methods lead us into exploring advanced image recognition techniques using machine learning.
- Section 5.3: Beyond the architecture of neural networks, the data itself is crucial. This section covers the collected datasets and deduces the reduced temperatures T/T_F of the thermometry dataset using the analysis introduced in Section 4.2.
- Section 5.4: Here, we detail the architecture and training of the supervised deep conventional neural network. After training, we deduce the superfluid critical temperature in the BCS-BCS crossover and analyse the network’s inner workings using the *DeepLift* library [200].
- Section 5.5: This section covers the architecture and training of the unsupervised deep convolutional autoencoder neural network. Following training, we analyse the data structure in the latent space and extract positions of trend changes. This allows us to map the superfluid critical temperature in the BEC-BCS crossover independently of the rapid ramp (RR) technique.
- Section 5.6: Finally, we summarise and conclude the results.

5.1 Detecting the superfluid phase transition

As discussed in Section 4.3, determining the superfluid phase boundary requires detecting the onset of condensation across the BEC-BCS crossover. In the strongly interacting regime, a clear signature of condensation, such as the bimodal momentum distribution seen in a BEC, is absent. Therefore, the rapid ramp (RR) technique has been established as the state-of-the-art method for detection of condensation. However, while the rapid ramp reliably identifies changes in condensation [166], the detection of the onset of condensation can be affected by the perturbative nature of the ramp. This incorporates a diminishing signal-to-noise ratio of the bimodal fit close to the phase transition, and possibly the detection of preformed pairs in a “pseudogap” state above the phase transition [51].

The first effect arises from the nature of the bimodal fit introduced in Eq. (4.13). Above the phase transition, the fit does not yield a zero condensate fraction, but instead captures deviations from a purely Gaussian thermal background, resulting in a noise-dominated, non-zero condensate fraction. For an experimental determination of the phase boundary in the strongly interacting regime, however, it is essential to have precise knowledge of the experimental conditions under which the condensate fraction vanishes.

Signatures of condensation

A more refined determination of the phase boundary requires techniques that can detect condensation beyond the RR. While TOF images do not provide a clear signature of condensation, pairing statistics still influence the momentum distribution in a condensed gas. However, as discussed in Section 4.3, the momentum distribution in TOF is influenced not only by pairing statistics (Fig. 2.6) but also by temperature, interactions, and the inhomogeneous trapping potential. This is evident when comparing the similar effects of finite temperature (Fig. 2.1) and pairing (Fig. 2.6) in a Fermi gas.

Nevertheless, in-situ images have been shown to exhibit small deviations from a non-interacting Fermi gas, caused by a finite condensate fraction [8]. Although no model function exists to directly infer condensation from these deviations, machine learning techniques can overcome this limitation by rendering a universal function approximator, trained to predict condensation from TOF images. By training a network on data spanning the strongly interacting regime, the network learns not only to detect the phase transition but to do so through generalised pattern recognition. To this end, we employ deep convolutional neural networks – an architecture proven to be ideal for image recognition tasks [178] – directly on TOF images.

5.2 Detecting phase transitions with machine learning techniques

Machine learning techniques have found diverse applications in physics [56, 57]. These range from simple optimisation tasks of experimental systems [59], such as the accelerated optimisation of experimental apparatuses by analysing the effect of a large parameter space [60], to improved detection methods, like absorption imaging with a single atom image [201]. Other applications include constructing quantum many-body state wavefunctions [202, 203] from experimental data.

In this thesis, the focus is on detecting phase transitions in experimental data using machine learning techniques. Neural networks have already been applied to theoretical data to extract phase transitions through supervised learning [204, 205]. Unsupervised neural networks have also been employed to detect phase transitions in theoretical data, utilising methods such as learning by “confusion” [61], principal component analysis (“PCA”) [61, 62, 64], and autoencoders [63, 64].

For experimental data, machine learning provides an ideal framework to detect phase transitions, as it can mask imperfections like imaging noise through the network architecture. Supervised machine learning has been used to identify phase transitions in topological systems [65], and later, unsupervised models have been applied to detect the phase transition in similar systems [66, 206, 207].

In this thesis, both a supervised neural network [67] and an autoencoder neural network [68] are applied to experimental TOF data to infer the phase transition.

5.3 Acquiring datasets across the phase transition

This section provides an overview of the recorded datasets used in this chapter. As previously established, the machine learning models in this thesis detect the phase transition from TOF images. As this does not directly yield the critical temperature T_C , the experimental settings corresponding to the phase transition are then repeated to deduce the temperature from corresponding in-situ images (refer to Section 4.2).

5.3.1 Controlled heating of a Fermi gas

The experimental settings at the determined phase transition are defined by the magnetic field, dipole trap power and corresponding ramps. To detect the phase transition, temperatures above and below the critical temperature T_C must be analysed for each magnetic field setting.

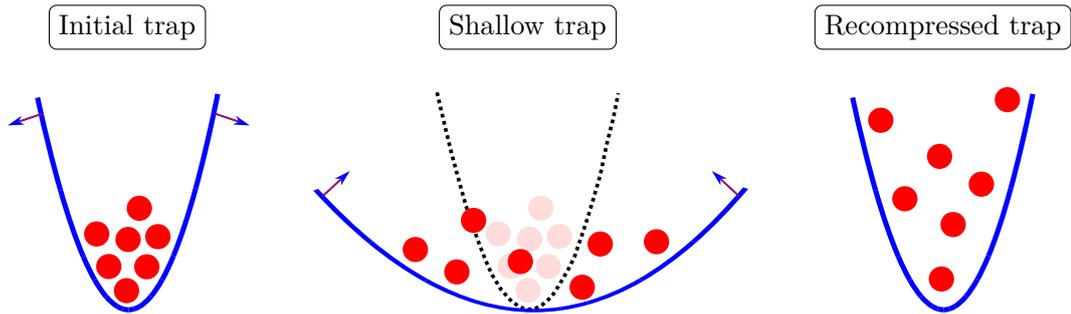


Figure 5.1: Concept of controlled heating. The sketch illustrates the principle of the controlled heating mechanism used to alter the temperature of the Fermi gas. Initially, the gas is thermalised in the initial trap. A sudden ramp down to a more shallow trap allows the atoms to expand. After a variable heating time t_{heat} , the trap is recompressed to the original configuration. Due to expansion in the shallow trap and subsequent recompression, the atoms gain potential energy, which is converted to temperature after additional thermalisation of 50 ms.

Heating is achieved through a controlled ramp of the dipole trap powers, as illustrated in Fig. 5.1 and described in [208, 209]. This process involves a sudden ramp down to a more shallow trap, followed by a recompression back to the initial trap setting after a variable heating time t_{heat} . Due to expansion in the more shallow trap and subsequent recompression, the atoms gain potential energy, which is converted to temperature after an additional thermalisation time of 50 ms in the recompressed trap.

During expansion, atom loss is minimised by keeping the horizontal dipole beam on, while the vertical dipole trap is completely turned off. Nevertheless, we find that the expansion is accompanied by strong atom loss, as indicated by the reduction in Fermi temperature T_{F} , shown in Fig. 5.2. However, this does not pose a problem, since the focus is on the temperature T normalised by the Fermi temperature of the gas T_{F} . A significant reduction in T_{F} consequently leads to an increase in the reduced temperature T/T_{F} .

5.3.2 Datasets

In total, three datasets with distinct purposes have been recorded, as outlined in the overview below. The overview includes information on the types of images captured and the magnetic fields at which data have been collected. The following sections provide further details about these datasets.

1. Training dataset

- Contains TOF images and CF from RR images
- Fields $B = 727.6 \text{ G}, 757.1 \text{ G}, 786.6 \text{ G}, 835.8 \text{ G}, 862.4 \text{ G}, 894.8 \text{ G}, 914.5 \text{ G}$
- ~ 7 repetitions per unique field and t_{heat} combination

2. Neural network dataset

- Contains TOF images only
- *BEC Dataset*: Fields $B = 723.7 \text{ G}, 728.6 \text{ G}, 733.5 \text{ G}, 739.4 \text{ G}, 744.4 \text{ G}, 751.2 \text{ G}, 759.1 \text{ G}$
- *BCS Dataset*: Fields $B = 759.1 \text{ G}, 775.8 \text{ G}, 803.4 \text{ G}, 835.8 \text{ G}, 855.5 \text{ G}, 880.1 \text{ G}, 914.5 \text{ G}$
- ~ 13 repetitions per unique field and t_{heat} combination

3. Thermometry dataset.

- Contains in-situ images via cameras Andor1 (top) and Andor2 (side) (check Section 4.1)
- *BEC Dataset*: Fields $B = 723.7 \text{ G}, 728.6 \text{ G}, 733.5 \text{ G}, 739.4 \text{ G}, 744.4 \text{ G}, 751.2 \text{ G}, 759.1 \text{ G}$
- *BCS Dataset*: Fields $B = 759.1 \text{ G}, 775.8 \text{ G}, 803.4 \text{ G}, 835.8 \text{ G}, 855.5 \text{ G}, 880.1 \text{ G}, 914.5 \text{ G}$
- ~ 30 repetitions per unique field and t_{heat} combination

1. Training dataset

For the supervised neural network in Section 5.4, labels of the corresponding condensate fraction CF must be recorded alongside TOF images. To achieve this, a training dataset is recorded for several magnetic fields around resonance, including temperatures below and

above the critical temperature. Each datapoint in this dataset comprises a TOF image and the condensate fraction derived from a corresponding rapid ramp image.

The TOF images are captured using standard absorption imaging at the field of preparation. To ensure a high signal-to-noise ratio while maintaining a mapping to momentum space, the expansion time for all TOF images is limited to 5 ms. Although this is shorter than the $T/4 = 15$ ms of the remaining harmonic potential, even a full development of the original momentum distribution after a $T/4$ expansion would be obscured by interactions. Hence, the neural network can still learn to extract relevant signatures from the density distribution after the given expansion time.

Details about the training process are covered in Section 5.4.1.

2. Neural network dataset

After training, an additional dataset is acquired to determine the phase transition. This dataset contains more magnetic field settings and does only contain TOF images, without supplementary information or labels, such as the condensate fraction.

The neural network dataset comprises datasets from two different dates, namely from 4th March 2022 and 12th April 2022. Former dataset comprises fields mostly on the BCS side of the crossover, namely at magnetic fields of $B = 759.1$ G, 775.8 G, 803.4 G, 835.8 G, 855.5 G, 880.1 G and 914.5 G, and will be referred to as the “BCS dataset”, henceforth. The latter dataset comprises magnetic fields closer the BEC limit, which are $B = 723.7$ G, 728.6 G, 733.5 G, 739.4 G, 744.4 G, 751.2 G and 759.1 G. This dataset will be called the “BEC dataset”.

Both datasets include ~ 13 repetitions for each unique combination of heating time t_{heat} and magnetic field B , while the parameters are randomly shuffled during recording. The following analysis of this dataset via the supervised neural network or the autoencoder network will yield a critical heating time τ_C , corresponding to the onset of condensation.

3. Thermometry dataset

Once the critical heating time τ_C is determined, the sample’s temperature and Fermi energy are measured to map the superfluid phase boundary. This is done in accordance with the thermometry described in Section 4.2.

To achieve this, an additional thermometry dataset has been recorded. This dataset contains in-situ density profiles and is recorded in an alternating fashion with the corresponding TOF images of the neural network dataset to mitigate the influence of long-time system drifts. The 3D density profile of the clouds in this dataset is then determined via the inverse Abel transformation, introduced in Section 4.2.1. To ensure low noise in the thermometry, the dataset includes ~ 30 repetitions. Additionally, a few in-situ images are taken from the side via Andor2 (see Section 4.1) to extract the aspect ratio for reconstructing the 3D density.

First, the 3D density profile is used to yield the temperature of the cloud via a fit of the virial expansion of the equation of state to the density profile, as explained in Section 4.2.3. The density of the gas also gauges the temperature scale to the energy of the system, that is the Fermi temperature T_F .

Condensation occurs below a critical temperature T_C/T_F . In our inhomogeneous trap, the density n_σ is highest in the trap centre, and hence the reduced temperature T/T_F is

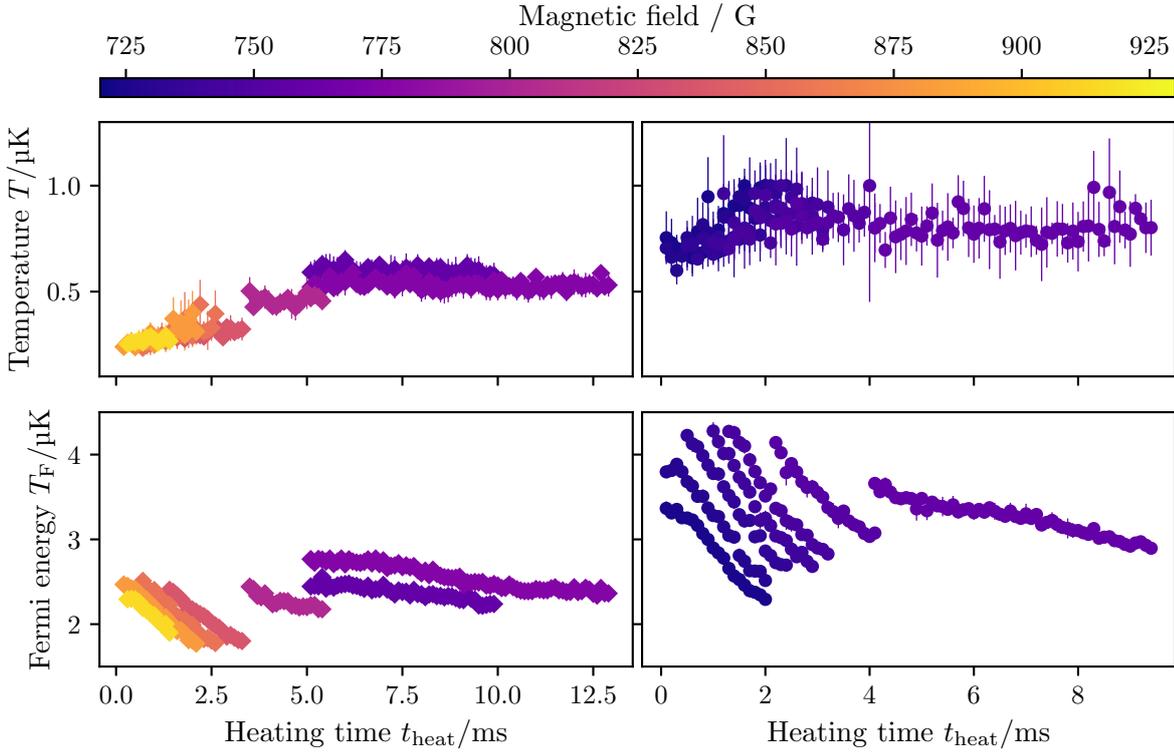


Figure 5.2: Thermometry results of the neural network dataset. The figure shows the thermometry results from the thermometry dataset calculated with the approach in Section 4.2. **Upper row:** Temperature T as a function of heating time t_{heat} . **Lower row:** Fermi temperature T_{F} as a function of heating time t_{heat} . The thermometry dataset comprises the BCS and BEC datasets as listed in Section 5.3.2, which are divided into the **left** and **right** column, respectively.

lowest here, with $T_{\text{F}} \propto n_{\sigma}^{2/3}$ (refer to Section 2.4). Consequently, condensation starts in the trap centre when the reduced temperature drops below the critical temperature $T_{\text{C}}/T_{\text{F}}$. Henceforth, the Fermi temperature T_{F} is defined with respect to the density in the trap centre.

The results of the thermometry are shown for the BEC and BCS datasets in Fig. 5.2. The figure illustrates that our controlled heating implementation primarily reduces the density and Fermi temperature, while the temperature of the gas stays rather constant. This yields an increase of the reduced temperature T/T_{F} , as anticipated.

It is worth noting that the thermometry dataset requires substantial repetitions and measurement time. To reduce this, sampling of the recorded heating times in the in-situ datasets has been concentrated around the critical heating time τ_{C} , as indicated by the RR measurements of vanishing condensate fraction.

To determine the critical temperature $T_{\text{C}}/T_{\text{F}}$ at specific magnetic fields, both the temperature T and Fermi temperature T_{F} are modelled by a linear fit. Around the critical heating time τ_{C} , changes in temperature and Fermi temperature are well approximated by a linear fit [134]. The exact value of both the critical temperature T_{C} and the Fermi temperature T_{F} at the critical heating time τ_{C} are deduced from best-fit parameters at τ_{C} . Furthermore,

the interaction parameter $1/(k_F a)$ is defined using the Fermi momentum $k_F = \sqrt{2mk_B T_F}/\hbar$. Note, that the standard error from the best-fit evaluation is reduced compared to the standard error shown in Fig. 5.2, since the linear fit incorporates numerous datapoints.

5.4 Supervised learning of the superfluid phase transition

The first approach in detecting the superfluid phase transition utilises a deep convolutional neural network trained to predict the condensate fraction CF from the momentum distribution in TOF images. The principles of both training and subsequent predictions of the CF via the neural network are sketched in Fig. 5.3.

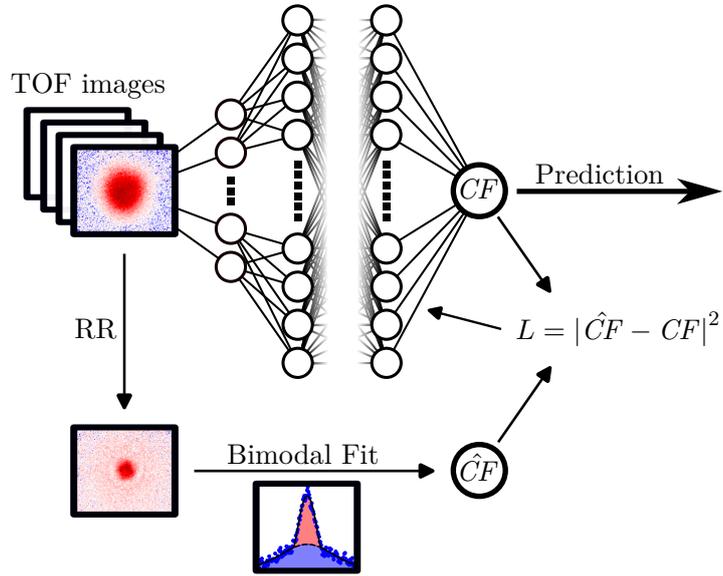


Figure 5.3: Supervised training of the neural network. The sketch illustrates how the convolutional neural network is trained to predict the condensate fraction CF from TOF images. To this end, a training dataset is prepared, comprising TOF images with corresponding labels of the condensate fraction \hat{CF} . The labels have been determined via a bimodal fit to RR images at the same experimental settings as the TOF images. During training, the weights of the network are adjusted to minimise the loss function $L = |\hat{CF} - CF|^2$. The weight adjustment is simplified here, for a more complete image refer to Section 4.4.2.

For training, the input TOF images have to be labelled with their corresponding condensate fraction \hat{CF} . A training dataset is prepared, assigning each TOF image with a corresponding \hat{CF} label. The TOF images are taken after a fixed expansion time of 5 ms without any magnetic field ramps after preparation. In an alternating fashion, the same sample is prepared but imaged after the rapid ramp (refer to Section 4.3), which yields the corresponding \hat{CF} . During training, the weights of the network are adjusted to minimise the loss function L , as explained in Section 4.4.2.

After training, the neural network is able to predict the condensate fraction – without additional labels – from new data that has not been seen during training [133, 134].

5.4.1 Training of the neural network

Before training starts, the weights of the neural network are randomly initiated. The training process adjusts the weights of the deep convolutional neural network to minimise the deviations from a desired output for a given input, expressed in the loss function L (refer to Section 4.4.2). Here, the inputs are TOF images, and the predicted output is the condensate fraction CF .

Training is conducted using the training dataset, introduced in Section 5.3.2. Since the goal is to enable the neural network to predict the condensate fraction from TOF images for arbitrary interactions $1/(k_F a)$, the training dataset includes various magnetic fields spanning the strongly interacting regime $-1 < 1/(k_F a) < 1$. Additionally, the training dataset covers data of different temperatures, i.e. heating times t_{heat} , encompassing both the superfluid and normal phase. The phase transition, in terms of the heating time t_{heat} , is estimated by the diminishing condensate fraction after the RR, which helps in selecting an appropriate sampling range for data collection.

In principle, the network could learn to predict correlation between atom loss during heating and the corresponding condensate fraction. However, this issue is mitigated by including data from many different magnetic fields in the dataset, each with different initial atom numbers. This forces the network to rely on more subtle features in the TOF images to predict the CF .

Network architecture

The network architecture resembles a deep convolutional neural network (refer to Section 4.4). This constitutes several convolutional layers with alternating stacks of max pooling layers, followed by several dense layers. The network architecture is illustrated in Fig. 5.4 and summarised in Table 5.1.

The input layer has a size corresponding to the size (225, 255) of the cropped and centred input TOF images. Directly after the input layer, two resizing layers adjust the input image size by upsampling by $\times 2$ and downsampling by $/3$ to (150, 170). The reason behind these layers will be explored during the explanation of “transfer learning and fine-tuning”, later in this section. After resizing, the output of size (150, 170) operates as the effective input to the neural network.

The deep convolutional neural network can be viewed as a pipeline for “feature extraction”, followed by the “classification” of these features to yield a prediction of the CF . The “feature extraction” is handled by several convolutional layers, which are each followed by a max pooling layer. As discussed in Section 4.4, these layers are tailored for image recognition.

After features have been extracted, they need to be classified to yield the prediction of the neural network. To this end, the output of the last max pooling layer is flattened and fed to several fully connected dense layers, which “classify” extracted features. Dropout layers between the dense layers are only relevant for training, during which randomly assigned

weights are cut off to prevent overfitting [181] with a dropout rate of 0.5. Finally, the output of the last layer is a single number, representing the predicted condensate fraction CF .

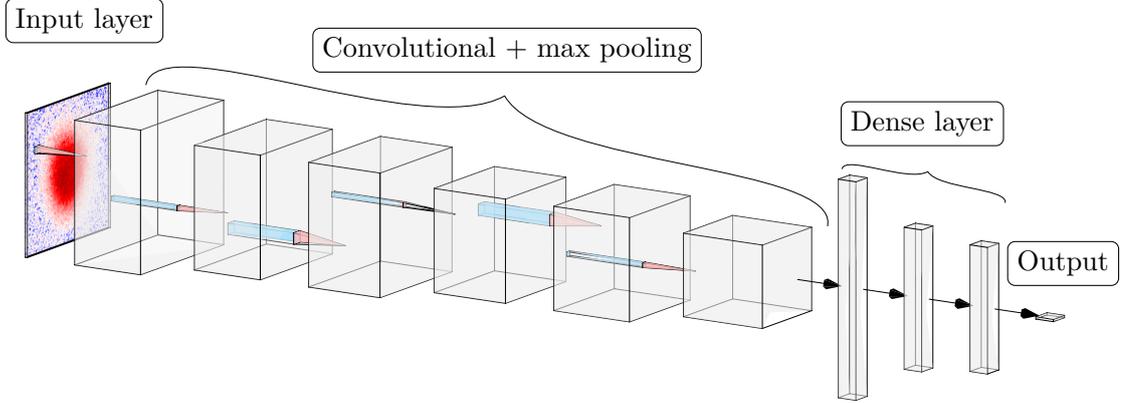


Figure 5.4: Architecture of the supervised neural network. The sketch showcases the architecture of the employed convolutional neural network in Section 5.4. Roughly speaking, the network is divided into alternating convolutional and max pooling layers with subsequent dense layers. Dropout layers between the dense layers are not shown for simplicity. The graphic has been created with [210]. Note that the layer sizes are not to scale.

Table 5.1: Architecture of the supervised deep convolutional neural network. This table presents the layer-by-layer architecture of the employed deep convolutional neural network [67], created using *tensorflow* [211]. The first two layers i and ii, simply rescale the image for the first functional layer 1. In total, the neural network has 11601591 parameters.

Layer index	Layer type	Output shape	Number of parameters
0	Input	(225, 255, 1)	0
i	Upsampling 2D	(450, 510, 1)	0
ii	Downsampling 2D	(150, 170, 1)	0
1	Convolutional 2D	(150, 170, 30)	300
2	Max pooling 2D	(75, 85, 30)	0
3	Convolutional 2D	(75, 85, 40)	30040
4	Max pooling 2D	(37, 42, 40)	0
5	Convolutional 2D	(37, 42, 50)	50050
6	Max pooling 2D	(18, 21, 50)	0
7	Flatten	(18900)	0
8	Dense	(600)	11340600
9	Dropout	(600)	0
10	Dense	(300)	180300
11	Dropout	(300)	0
12	Dense	(1)	301

Transfer learning and fine-tuning

During the works on this project, the experimental apparatus was subject to slight changes and adjustments. These changes comprise a change to the optical dipole trap as well as an adjustment of the imaging system to compensate for a camera replacement (Andor0, refer to Section 4.1). Unfortunately, after these changes the already trained network from [133, 134] no longer produced accurate and reasonable predictions of the condensate fraction anymore. To restore its predictions' accuracy, the network had to be retrained using images from the updated system.

Although the old network [133, 134] was no longer effective, it had already been trained on a large dataset and has successfully learned to extract signatures from TOF images to predict the condensate fraction. This “knowledge”, stored in the network’s weights, can be leveraged to retrain the network on a smaller (new) dataset. The retraining process is known as “transfer learning” and is commonly used in many scenarios such as image classification, where a large pre-trained neural network serves as a starting point to train a network on a smaller dataset of new data [197–199].

The network architecture for the modified system is identical to the architecture of the previous system [133, 134], but adds an upsampling $\times 2$ and downsampling $/3$ layer after the input layer to account for the changed pixel size ($16\ \mu\text{m} \rightarrow 13\ \mu\text{m}$) and magnification ($3.42 \rightarrow 4.11$) of the new camera. As shown in Table 5.1, these layers have no parameters and are not adjusted during training. Instead, they allow for having the input size (225, 255) of the new system to be mapped onto the input size (150, 170) of the previous architecture, preserving the relative size of features in the input images.

Preparing the training dataset

Before any modifications to the experimental apparatus, the network was trained on a dataset of 7895 labelled data, consisting of the fields listed in Section 5.3.2 [133, 134]. For transfer learning, a new training dataset with the modified system has been recorded, comprising a smaller dataset of 2245 datapoints. This dataset is randomly split into a training and validation set, comprising 90 % and 10 % of the full dataset, respectively. The training set is used for training, while the validation set is reserved for testing the network’s performance on unseen data.

We train the neural network with stochastic gradient descent using the Adam optimiser with a learning rate of 1.5×10^{-4} [192] and a batch size of 30 [67].

As discussed in Section 5.1, the idea of the network is to predict the onset of condensation at the critical heating time t_C , where the rapid ramp produces unreliable data. To address this, the training set is masked and modified as shown in Fig. 5.5. The figure shows the labels from the RR, as well as predictions of the network after training. For large condensate fractions, the rapid ramp can be trusted, and hence data with these labels are used for training. However, below a small but finite condensate fraction of 3 % – 5 %, the rapid ramp becomes less reliable, and the measured CF values are dominated by noise in the bimodal fit, as explained in Section 4.3. The cut-off threshold is chosen for each magnetic field in the training set separately, considering the data range and noise of the recorded CF values. Data from the training set, which have labels below the cut-off threshold value are removed for

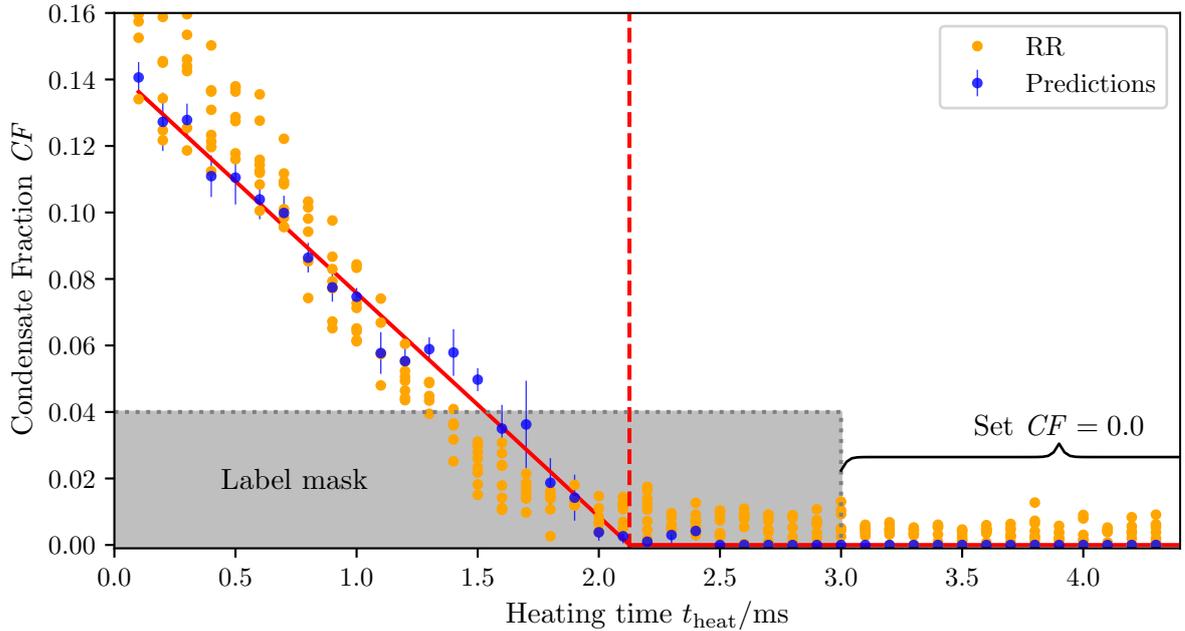


Figure 5.5: Labels and predictions of the training dataset. The training dataset contains TOF images as well as the condensate fraction CF from the RR as labels (orange points). Below a threshold, the labels from the rapid ramp can not be trusted to yield a sharp phase transition (grey area, label mask). Training data below this threshold are in the “label mask” and are removed for training. Beyond a large heating time, the condensate fraction labels can be assumed to be $CF = 0$ and are set accordingly. After training, the network produces predictions (blue points), which continue a sharp trend of decreasing CF towards $CF = 0$. The critical heating time (red dashed line) is yielded from a linear-constant fit to the predictions as given in Eq. (5.1).

training. At very large heating times, the CF labels can safely be assumed to be zero, and these data points are included in the training set with CF set strictly to zero.

Training and validation of the neural network

To accelerate the training process, the technique of transfer learning is employed. Instead of initialising the network with random weights, the weights from the old network [133, 134] are used for the convolutional layers, i.e. layer 1 – 6 in Table 5.1. The weights in the subsequent classifier part of the network, which consists of the fully connected layers 7 – 12, are randomly initiated.

A key metric to judge if the training has converged is the loss after each training epoch [181]. The loss after each training epoch is shown in Fig. 5.6 for the training and validation set, respectively. For the first 15 epochs, the convolutional layers are frozen to allow for transfer learning. This means, that the weights of the convolutional layers are not adjusted, while all unfrozen layer are adjusted via stochastic gradient descent, as described in Section 4.4.2. In this way, the network adopts the old weights of feature extraction in the convolutional layers by only training a new classifier.

After 15 epochs, the loss of both the training and validation set converge, and fine-tuning

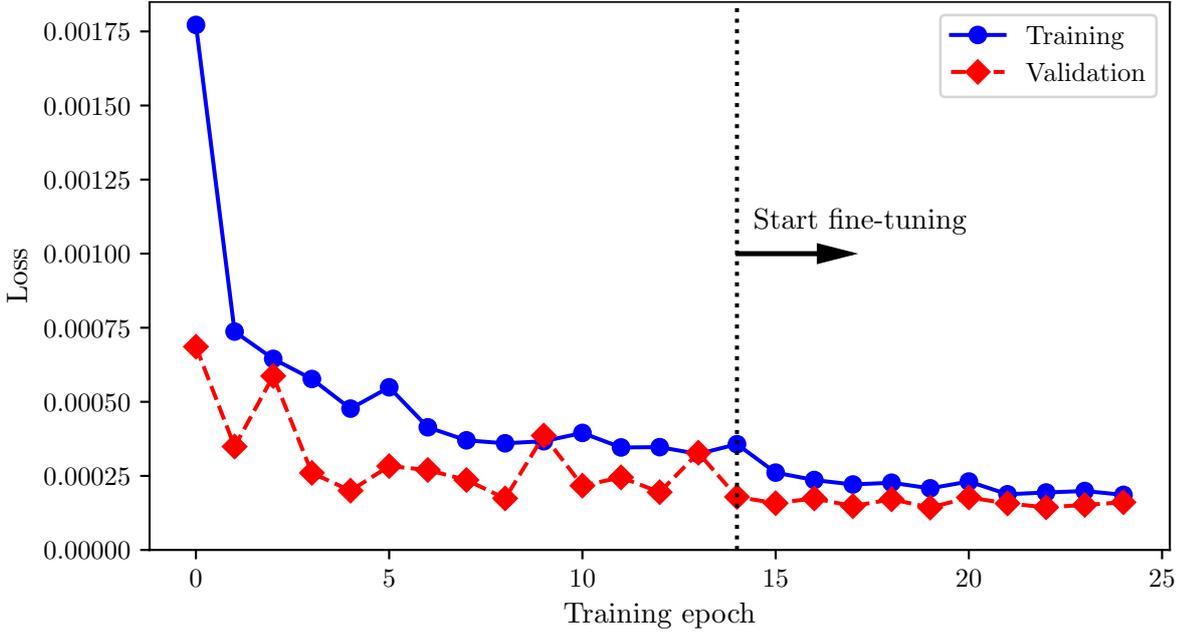


Figure 5.6: Training and validation loss during transfer learning. This plots shows the loss (refer to Section 4.4.2) over many training epochs of the training and validation set, which are randomly assigned datapoints in a ratio of 9 : 1 of the training dataset. For the first 15 epochs, the network is trained via transfer learning using the weights of the old network [133, 134] for the convolutional layers. For the next 10 epochs, the network is fine tuned by opening all layers to adjustments during gradient descent.

of the network is initiated by unfreezing all layers. This adjustment enables the network to refine the weights of the convolutional layers, adapting them to the modified system. During this phase, the learning rate of the Adam optimiser [192] is reduced from 1.5×10^{-4} to 1×10^{-5} . After another 10 epochs of fine-tuning, the training completed as both the training and validation losses converge to a lower value.

Upon completing training, the network can accurately predict the condensate fraction from TOF images, as demonstrated in Fig. 5.5.

Generalisation of validation

After training, it is essential to verify whether the network has genuinely learned a generalized method of feature extraction to predict the condensate fraction CF from the momentum distribution in TOF images, rather than merely learning correlations with obvious parameters such as the summed optical density (i.e. the atom number) or the cloud size.

If the network only learns these correlations, it would be possible to train a similarly performing network using Fermi gas fits [8] instead of the full TOF image. Indeed, we observe that a network can be trained on best-fits derived from Fermi gas fits [134], however, leading to a higher general loss compared to a network trained with the full TOF images.

As an additional test, we have employed the neural network on a $|13\rangle$ mixture, even though the network has been trained on a $|12\rangle$ mixture. Remarkably, the network still

produces predictions close to the results obtained with the RR [134], even though it had never encountered images of this mixture before.

As we will see in Section 5.4.2, the network is also able to predict the CF for fields which have not been part of the training dataset. This suggests that the network has indeed learned a general approach in extracting the condensate fraction from the momentum distribution in TOF images. In section Section 5.4.4, this will be illustrated by visualising the importance of regions in the input activation map of TOF images on the predicted CF across the crossover.

5.4.2 Determining the phase transition

With the neural network trained to predict the condensate fraction directly from the momentum distribution in TOF images, the next step is to determine the critical heating time τ_C at which the onset of condensation occurs. To this end, the trained neural network is employed to predict the CF from TOF images in the neural network dataset (refer to Section 5.3.2). For each unique combination of the magnetic field and heating time t_{heat} , the TOF images have been measured for ~ 13 repetitions. Next, the predicted CF are averaged for the ~ 13 predictions. The mean of the predicted CF is used henceforth.

To determine the critical heating time τ_C , the predicted condensate fraction CF is modelled by a linear-constant fit as already shown in Fig. 5.5. The model function reads

$$CF(t_{\text{heat}}) = \begin{cases} a \cdot (t_{\text{heat}} - \tau_C) & \text{if } t_{\text{heat}} \leq \tau_C \\ 0 & \text{else,} \end{cases} \quad (5.1)$$

with a being a fit variable and τ_C being the critical heating time.

In Fig. 5.7, the fits are displayed along with the predicted condensate fraction for exemplary magnetic fields of the dataset. At magnetic fields close to unitarity, the condensate fraction at zero heating $t_{\text{heat}} = 0$ deviates strongly from the linear behaviour observed for larger heating times. Therefore, the fitting range for these fields is limited to the linear regime.

By extracting the critical heating time τ_C , the superfluid phase boundary is determined in terms of the heating time t_{heat} . An overview of all determined critical heating times for all considered magnetic fields is presented in Fig. 5.8.

Finally, the determined critical heating times are mapped to the thermometric values of the samples by using the analysis of the thermometry dataset in Fig. 5.2. This yields the critical temperature T_C as well as the Fermi temperature T_F and Fermi momentum $k_F = \sqrt{2mk_B T_F}/\hbar$. Together, these values describe the envelope of the superfluid phase boundary in terms of T_C/T_F and $1/(k_F a)$, as shown in Fig. 5.9.

5.4.3 Critical temperature

The superfluid phase diagram in Fig. 5.9 depicts the determined superfluid critical temperature in a homogeneous Fermi gas due to the definition of the Fermi gas within LDA, using the density at the centre of the trap. In summary, the resulting phase diagram shows the critical temperature for interactions strengths from $1/(k_F a) \approx -0.4$ to $1/(k_F a) \approx 1.0$. From the BCS limit up to an interaction parameter of $1/(k_F a) \approx 0.5$ the critical temperature increases. For weaker interactions approaching the BEC limit, the critical temperature remains roughly

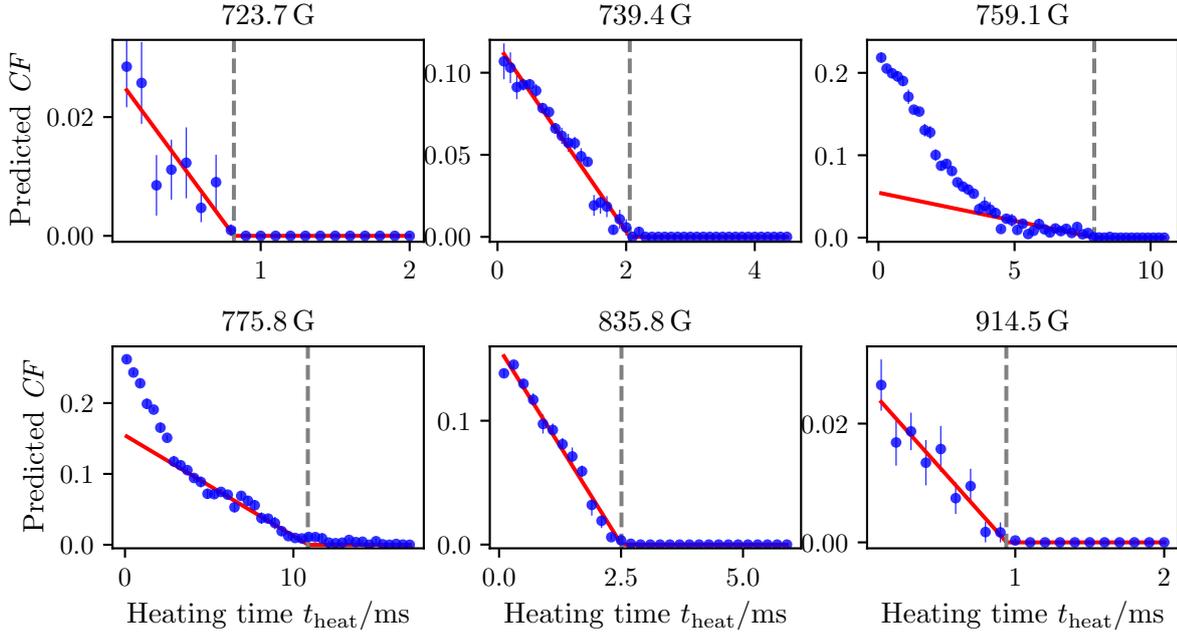


Figure 5.7: Determining the critical heating time. The figure shows the predicted CF for six exemplary fields from data in the neural network dataset (refer to Section 5.3.2). For each heating time t_{heat} , ~ 13 repetitions have been measured, and the datapoints (blue) show the averaged prediction while the error corresponds to the standard error of the mean. A linear-constant function (red, refer to Eq. (5.1)) is fitted to the data to yield the critical heating time τ_{heat} (grey, dashed line). Note that the **upper row** corresponds to the “BEC dataset”, while the **lower row** corresponds to the “BCS dataset”, defined in Section 5.3.2.

constant, or possibly, indicates a weak decline towards the BEC limit. This presents the first experimental confirmation of an increasing critical temperature from the BCS limit, extending beyond unitarity and approaching the BEC limit.

Our results show a striking resemblance with the extended GMB theory presented in [19]. Towards the BCS limit, our results align more closely with the t-matrix approach findings in [45]. Overall, our results are also consistent with results from quantum Monte Carlo simulation reported in [128]. Due to the uncertainty in our data, a peak in the critical temperature at an interaction parameter of $1/(k_{\text{F}}a) \approx 0.6$ [19] is only suggested but can not be conclusively confirmed.

In conclusion, we have demonstrated that the superfluid phase transition can indeed be detected via enhanced image recognition using supervised machine learning techniques. This measurement of the critical temperature presents the first precession measurement of the superfluid phase boundary in the strongly interacting regime for various interaction parameters.

5.4.4 DeepLift - Insights into the network

The final question to address is: what does the neural network actually learn? As previously discussed, the neural network is able to predict the CF from data taken at interaction

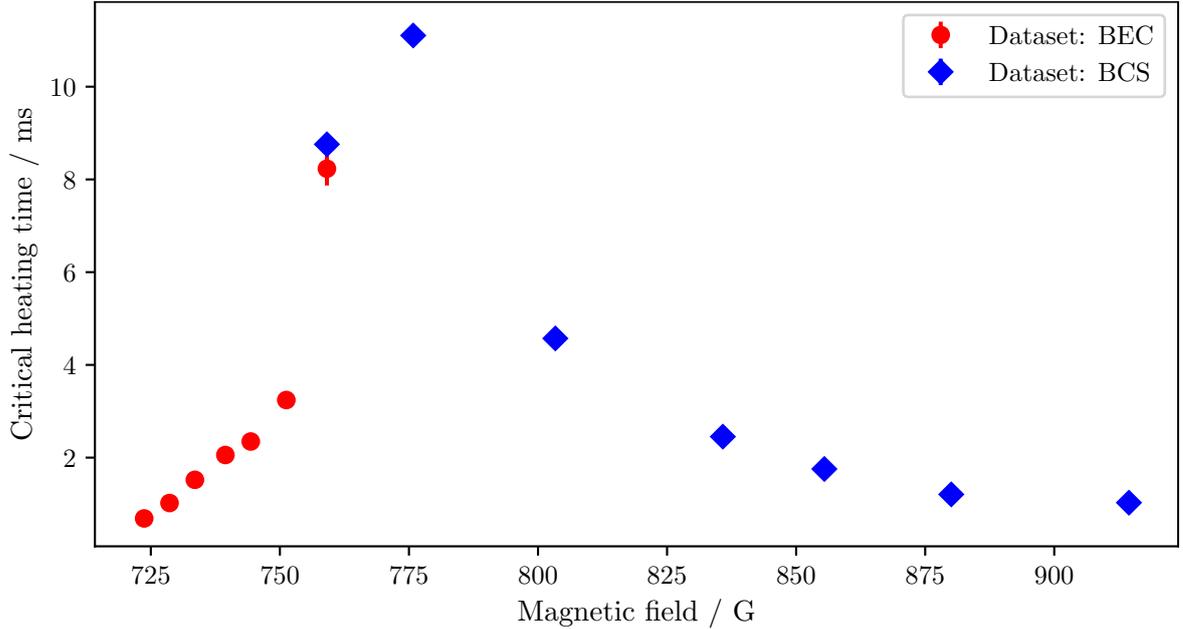


Figure 5.8: Overview of the determined critical heating times. The plot shows the determined critical heating times τ_C for all magnetic fields included in the neural network dataset (refer to Section 5.3.2). The dataset comprises two different logs, which have been captured on different days, labelled as the “BEC” and “BCS” datasets, respectively, describing which limit of the crossover they approach. Errorbars denote the standard error and are mostly covered by the size of the datapoints.

parameters, i.e. magnetic fields, that have not been seen during training. This suggests that the neural network is indeed able to generalise the detection of features in the momentum distribution from TOF images to determine the CF . However, it remains unclear which features the neural network focuses on to extract the CF .

Neural networks are often viewed as “black boxes” that produce a desired output for a given input, without transparency regarding their internal processes. Understanding the mapping from input to output would allow for a more interpretative understanding of how the neural network functions [212]. To gain these insights, we employ the DeepLift [200] algorithm on the trained network using the training dataset.

DeepLift assigns an importance score to the input neurons concerning a selected output neuron by using a backpropagation approach. This score quantifies the contribution of each input neuron to the output neuron’s value, considering all intermediate layers. In our case, the output neuron is chosen to be the final output neuron, yielding the CF . The input neurons are the pixels of the corresponding input images, showing the momentum distribution after TOF. These are referred to as the “activation map” of the input layer. Consequently, DeepLift generates an importance score “image” that has a one-to-one correspondence with the pixels of the input images from TOF.

Given the TOF duration of 5 ms, each pixel in the input image can be associated with a momentum k/k_F , assuming a non-interacting Fermi gas. The importance scores are radially averaged along momentum bins, mapping the importance scores to momenta. This analysis is conducted only for images at zero heating time $t_{\text{heat}} = 0$, since we aim to understand what

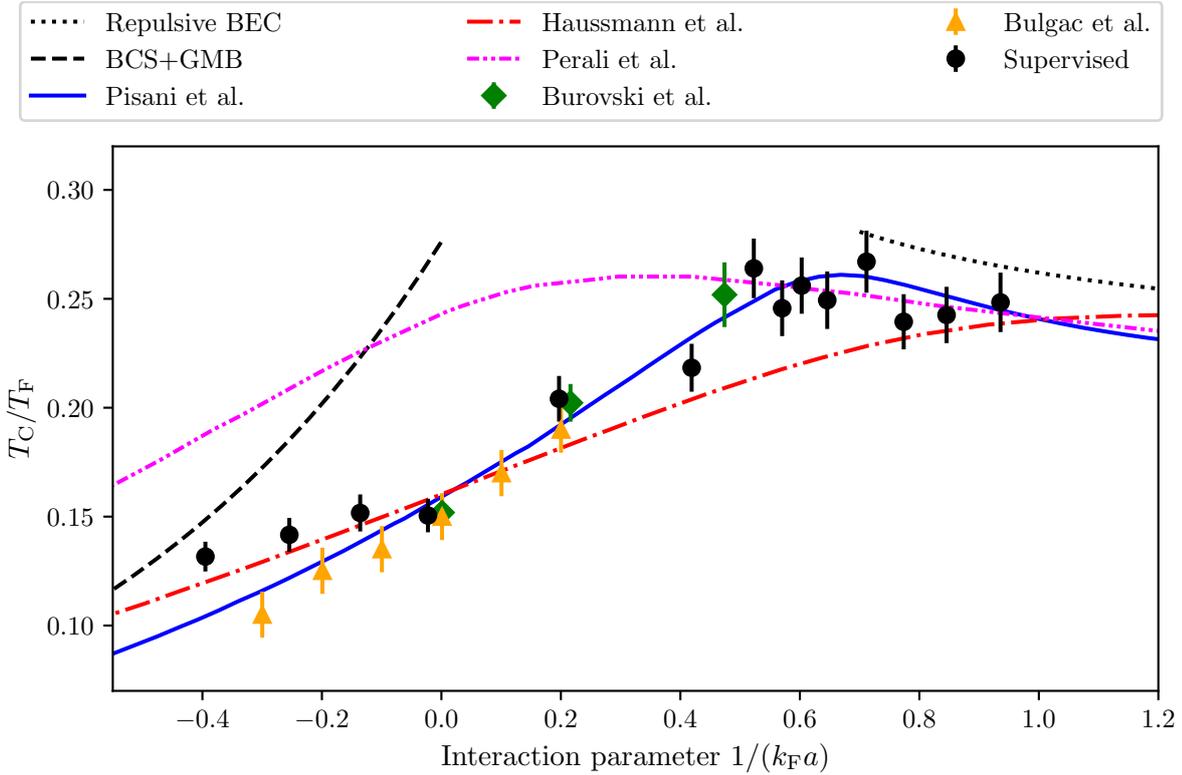


Figure 5.9: Critical Temperature in the BEC-BCS crossover via the supervised neural network. The figure presents the results of the superfluid critical temperature, which were determined via the supervised deep convolutional neural network (circles), adapted from [67]. The error comprises the standard error as well as a dominating 5% uncertainty due to anharmonicities of the trapping potential (refer to Section 4.2.1). Several theoretical calculations of the critical temperature T_C/T_F are shown alongside the results. In the limits, the analytical results for the weakly interacting BEC (Eq. (2.13)) (dotted line) and the exponentially decaying BCS result with GMB correction (Eq. (2.25)) (dashed line) are shown. Across unitarity an extended GMB correction [19] (solid line) and a self-consistent t-matrix approach [45] (dashed-dotted line) are plotted. Additionally, two quantum Monte Carlo results (diamonds [128] and upright triangles [47]) are presented.

features the network considers crucial when predicting the condensate fraction CF .

The importance scores for different interaction parameters $1/(k_F a)$ are shown in Fig. 5.10 as a function of the momentum bins k/k_F of the TOF images. As anticipated, the importance scores are higher in the centre of the TOF images on the BEC side, corresponding to low momenta. Moving towards unitarity and the BCS side, the importance scores shift towards higher momenta, reflecting Cooper pairing at higher momenta (refer to Fig. 2.6). This trend aligns with the evolution of pairing effects on the momentum distribution across the crossover, as illustrated in Fig. 2.6.

Momenta as high as the Fermi momentum $k = k_F$ are not relevant in the activation map due to the significant drop in the signal-to-noise ratio in the outer wings of the TOF images. Additionally, the TOF images represent a 2D projection of the 3D Fermi sphere, which smears out the Fermi edge. Since TOF images are taken after 5 ms expansion – less than

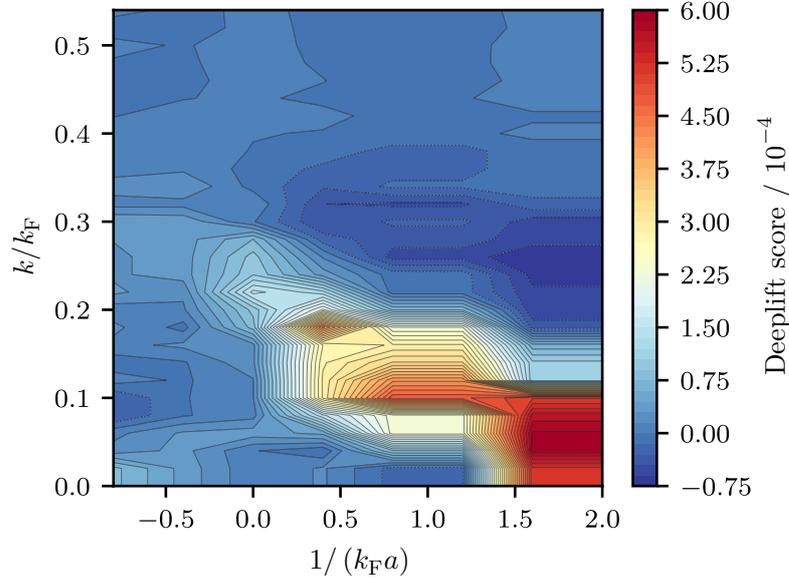


Figure 5.10: Importance scores of momentum bins of the input image for the prediction of the network output. The importance scores of pixels in the input image are determined via the *DeepLIFT* library [200]. To this end, only the coldest samples without any heating time from the training dataset are analysed. Pixels are correlated with their momentum after 5 ms of TOF, omitting interactions. Both axes are normalised with the Fermi momentum k_F , corresponding to a harmonic trap with N atoms. The image is taken from [67].

$T/4 \sim 15$ ms – the initial density also influences the momentum distribution. Lastly, the momenta are calculated assuming a non-interacting Fermi gas, but attractive interactions could slow down the expansion, leading to an underestimation of the calculated momenta.

Overall, the DeepLift approach shows that the highest importance scores indeed shift towards higher momenta when moving from the BEC limit towards unitarity and the BCS side. This indicates that the neural network effectively “learns” physical features from the momentum distribution in TOF data.

5.5 Autoencoder - Unsupervised detection of the superfluid phase transition

In the previous section, the superfluid phase transition in the strongly interacting regime has been determined using a supervised deep convolutional neural network. This section explores an alternative approach, utilising a different network architecture to identify the phase transition without relying on the rapid ramp technique for the creation of labels. Here, we employ an unsupervised autoencoder deep convolutional neural network, known as an autoencoder neural network.

The principle of the autoencoder neural network is illustrated in Fig. 5.11. As an unsupervised technique, its approach significantly deviates from the supervised deep neural network in

Section 5.4. The key idea behind the autoencoder is to reduce the dimensionality of the input data space, e.g., images, by constructing a unity function. Ultimately, the autoencoder is optimised to map input images to an output of the same size and dimensionality – essentially functioning as a unity function. However, by constraining the size and dimensionality of a hidden layer within the network, the autoencoder network is forced to compress the input images through this “bottleneck”, while optimising the subsequent layers to reconstruct the original image. The specific hidden layer is referred to as the “bottleneck layer”.

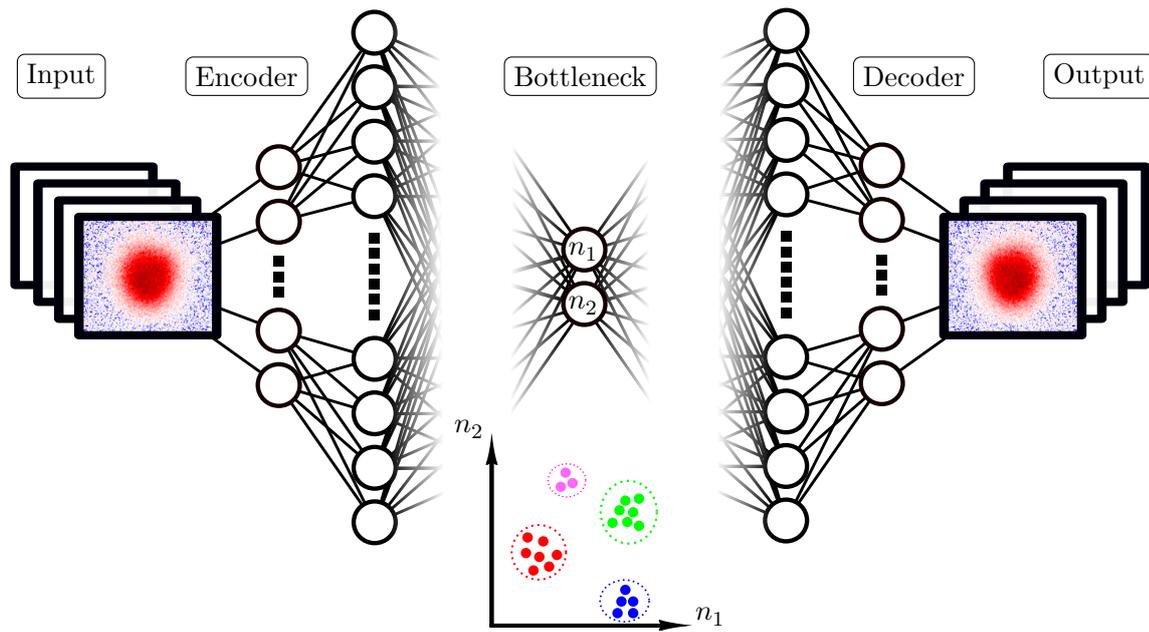


Figure 5.11: Principle of the autoencoder network. The sketch shows the working principle of the autoencoder neural network. Input and output have the same dimensionality and size. First, the input image is compressed by the encoder part of the network through the bottleneck layer. The bottleneck layer is made up of a small amount of neurons, here two neurons n_1 and n_2 . After the bottleneck layer, the network is optimised to give an output that reconstructs the input through the decoder part. After training, the distribution of input images in the latent space is analysed, which is spanned by n_1 and n_2 . Note that every datapoint in the latent space corresponds to one input image.

In this architecture, the layers preceding the bottleneck encode the input data into the low-dimensional space of the bottleneck layer, forming the “encoder”. Similarly, the layers following the bottleneck constitute the “decoder”.

For this thesis, the bottleneck layer consists of just 2 neurons, though configurations with more neurons are tested later in Section 5.5.2. To optimise the network to map the unity function, the weights are adjusted through supervised gradient descent. However, unlike for the supervised learning in Section 5.4, no additional labels are created here. Instead, the input images serve as input and labels at the same time, training the network to approximate a unity function using stochastic gradient descent (refer to Section 4.4.2).

After the autoencoder is successfully trained, the network maps input images onto them-

selves with minimal deviation. While this process alone does not reveal any novel structures, the low-dimensional representation that has been learned during training reveals structure in a complex representation of the input space [179]. To extract the low-dimensional representation from the bottleneck layer, the decoder part is stripped from the network, leaving the encoder as a standalone network with the bottleneck layers as its output. The coordinate space spanned by the outputs of the bottleneck layer is known as the “latent space”.

For distinctly different input images, the network organises the latent space such that distinct classes of input images cluster into separated regions [179]. Here, we train the autoencoder on TOF images that involve a second-order phase transition to a superfluid. While this suggests that no clear clustering is to be expected, a trend from superfluid to normal phase images may emerge, allowing for detection within the latent space.

5.5.1 Architecture of the autoencoder network

The architecture of the employed autoencoder neural network is detailed in Table 5.2. In principle, the architecture is similar to the supervised deep convolutional neural network in Table 5.1, but it features two convolutional neural networks that are symmetric around the central bottleneck layer.

The encoder begins with two convolutional layers, each being followed by a max pooling layer. Afterwards, the output of the last convolutional layer is flattened and compressed to the bottleneck layer via multiple fully connected dense layers. The decoder, which approximately mirrors the encoder, follows the bottleneck layer. After passing through some fully connected dense layers, the one-dimensional output is reshaped to a two-dimensional layer, which is subsequently processed by multiple transposed convolutional layers [181]. Finally, the output from the last convolutional layer matches the size and dimensionality of the input layer.

Subtleties of the architecture

After training, the output of the autoencoder network is expected to resemble the input as close as possible, with small deviations depending on the network architecture. Although, the network has been tested to be stable against changes in most hyperparameters of the given layers [133, 134], certain parameters and settings require careful adjustments to ensure that the network effectively learns the unity function. Key issues that have been identified during this process are summarised in Fig. 5.12.

An obvious issue arises if the activation function Φ (refer to Section 4.4) is not chosen properly. In many hidden layers, the data range does not influence the performance of the network, and hence a sigmoid activation function can be employed. However, in layers where the output range is critical, it is essential to use an activation function that covers the full relevant data range. Improper selection can lead to the output range, such as the optical density (OD) in TOF images, to be capped, resulting in the loss of important features. To avoid this, the whole network employs only ReLu activation functions, except for the bottleneck layer, which employs a purely linear activation function.

Caution has to be taken with the hyperparameters of the transposed convolutional layers (refer to Section 4.4.1). Transposed convolutional layers are conceptually the inverse of convolutional layers, intended to reconstruct the input from a known output. However,

Table 5.2: Architecture of the autoencoder neural network. The table presents the layer-by-layer architecture of the employed autoencoder neural network [68], created using *tensorflow* [211]. For clarity, the bottleneck layer 8 is highlighted. After training, the encoder part, comprising layers 0 to 8, is extracted, and the output of layer 8 constitutes the latent space. In total, the neural network has 17400963 parameters.

Layer index	Layer type	Output shape	Number of parameters
0	Input	(192, 192, 1)	0
1	Convolutional 2D	(192, 192, 32)	320
2	Max pooling 2D	(48, 48, 32)	0
3	Convolutional 2D	(48, 48, 64)	8256
4	Max pooling 2D	(16, 16, 64)	0
5	Flatten	(16384)	0
6	Dense	(512)	8389120
7	Dense	(512)	262656
8	Dense - Bottleneck	(2)	1026
9	Dense	(512)	1536
10	Dense	(512)	262656
11	Dense	(16384)	8404992
12	Reshape	(16, 16, 64)	0
13	Transposed convolutional 2D	(48, 48, 64)	36928
14	Batch normalisation	(48, 48, 64)	256
15	Transposed convolutional 2D	(192, 192, 32)	32800
16	Batch normalisation	(192, 192, 32)	128
17	Convolutional 2D	(192, 192, 1)	289

the operation can present issues, particularly related to the strides parameter. During the convolutional operation, the kernel moves along the input space in steps defined by the strides parameter. If the strides parameter is chosen bigger than the kernel size, not all input pixels contribute, leading to a grid-like pattern in the output – a critical issue in transposed convolutional layers if not properly managed.

To prevent these issues, the output images are always carefully checked after training, along with the training and validation losses, to ensure the network functions correctly.

5.5.2 Latent space - Detecting the phase transition in a low-dimensional representation

After setting up the network architecture, the autoencoder is trained on the neural network dataset (refer to Section 5.3.2). This dataset contains 5031 datapoints, with 90 % used for training and 10 % for validation. The covered heating times t_{heat} in this dataset range from no heating – corresponding to a large condensate – to high heating times that surpass the phase transition, as illustrated in Fig. 5.7. During training, the dataset is randomly shuffled and divided into batch sizes of 20. The autoencoder is trained over 15 epochs using the Adam optimizer [192] with a learning rate of 4×10^{-4} , after which the loss converges, as shown in

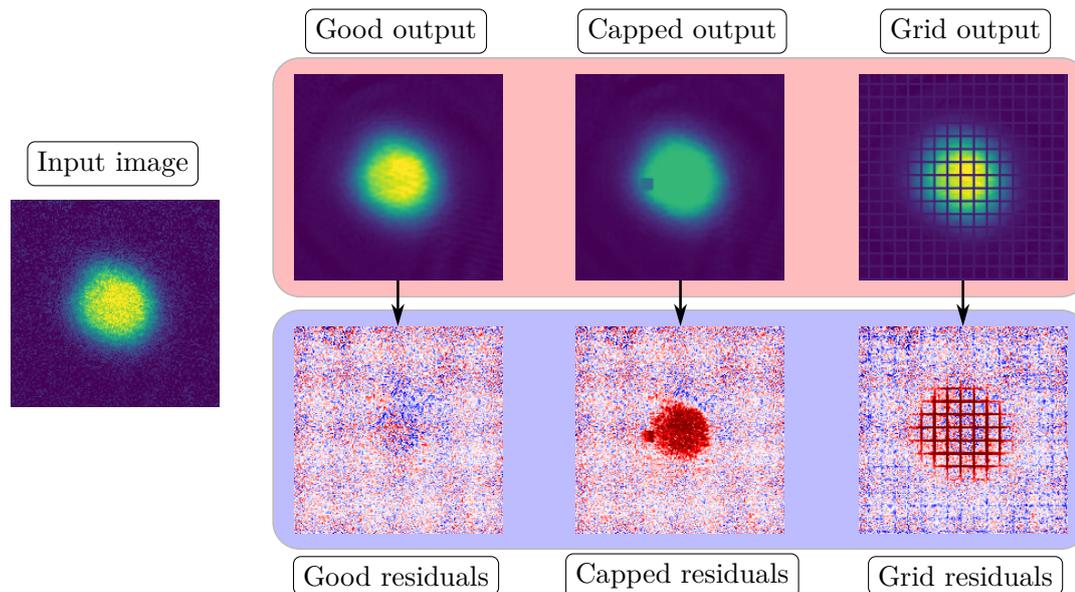


Figure 5.12: Output and residuals of the autoencoder network. The figures shows an output image along with its residuals with respect to its input image for a proper and two problematic network architectures. The first image shows the “good output” if the network has no major issues and the output resembles the input image as close as possible. On the contrary, the network is not able to reproduce the input image, if the activation functions are not chosen properly (“capped output”) or if the strides parameter of the transposed convolutional layers do not fit the kernel size (“grid output”). All output images and residuals have the same colourscale, respectively.

Fig. 5.19.

Following training, the decoder part is stripped off the autoencoder, and the encoder is used as a standalone network. Now, the output from the encoder provides a low-dimensional representation of the input dataset through the bottleneck layer. The coordinate system spanned by the two bottleneck neurons is the latent space.

Fig. 5.13 provides an overview of the latent space for all input images, where colour denotes the magnetic field (left) or the heating time t_{heat} (right). In the latent space, each TOF image corresponds to one datapoint, defined by the coordinate (Neuron1, Neuron2). Interestingly, the autoencoder organises images in the latent space along lines associated with distinct magnetic fields, which correlate with the interaction parameter. Along these lines, the images are sorted by heating time t_{heat} , which serves as a measure of the temperature of the gas (refer to Fig. 5.2). Remarkably, the autoencoder is not provided with any information about these quantities; it learns to infer these properties of the samples purely from the structure in the TOF images.

The lines for different magnetic fields appear to converge to a small region in the upper left corner of the latent space as the temperature increases. For lower temperatures, the images are more widely dispersed across the latent space, varying in location by magnetic fields. This behaviour is expected, as the latent space tends to cluster similar images, such as hot Fermi gas, closer together [8]. In other words, the autoencoder perceives minimal differences in TOF images with respect to interactions when the gases are hot and devoid of

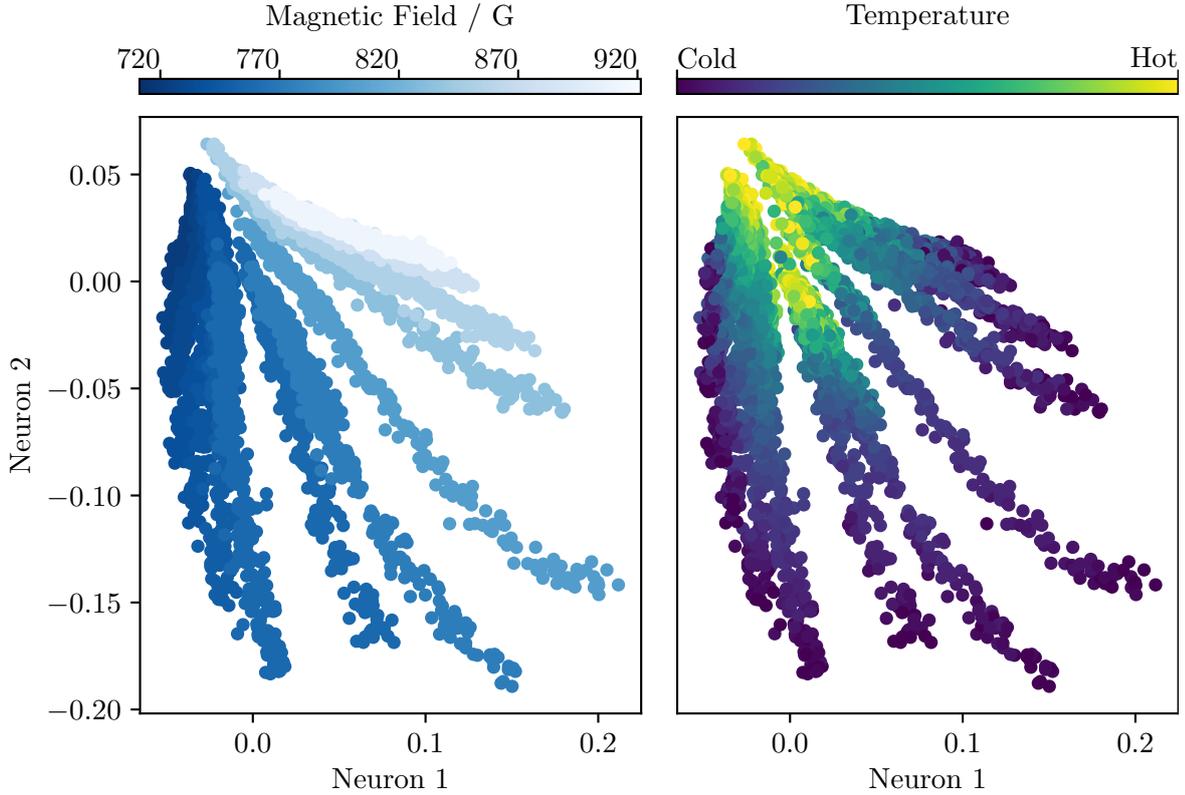


Figure 5.13: Latent space representation of the neural network dataset. The figure showcases the latent space distribution of the TOF images from the neural network dataset. The latent space is spanned by Neuron 1 and Neuron 2 from the bottleneck layer. Each datapoint corresponds to one TOF image. **Left:** The colour highlights the magnetic field of the corresponding TOF images. **Right:** The colour denotes the heating time t_{heat} of the TOF images. The temperature scale is qualitative, since every magnetic field series has a different range of heating times. Note that the autoencoder neither has information about the magnetic field or the heating time, but learns to sort the data in the latent space by nothing but the momentum distribution in TOF. Figure adapted from [68].

condensation. Cold gases with finite condensation, in contrast, exhibit more variation across different interactions, occupying a larger region in the latent space.

This behaviour alone is intriguing, as it suggests that the autoencoder can distinguish between temperature and interactions solely based on TOF images. A more quantitative discussion about the second-order superfluid phase transition, however, requires a deeper understanding of how data is arranged in the latent space. As demonstrated with the supervised neural network in Fig. 5.7, the phase transition manifests as a sharp onset of condensation at a specific heating time. Features in the arrangement within the latent space are expected to be more subtle, since the second-order phase transition does not directly alter the density distribution in TOF images but rather emerges as a discontinuity in derivatives of the energy, e.g., the heat capacity or the compressibility [165]. To identify features corresponding to the second-order phase transition, subtle trend changes in the arrangement of data within the latent space must be examined.

Finding trend changes in the latent space

Following a single magnetic field line from hot to cold samples in Fig. 5.13 suggests that the line is not described by a single straight line but rather follows a more complex curve. Especially, strong bending is observed for the coldest samples. However, the bending at the smallest temperatures can not be attributed to the phase transition, as the supervised analysis in Fig. 5.7 clearly indicates a substantial remaining condensate at these temperatures. Similarly, the hottest samples might exhibit additional structure in the latent space, although this is less pronounced due to the smaller area in the latent space occupied by hot gases. The most critical area to be examined is centred around the critical heating time that has been determined via the supervised analysis, omitting the coldest and hottest samples. Further details regarding the trimming of data in the latent space is presented in Fig. 5.15.

To identify trend changes, the coordinates in the latent space are averaged for the ~ 13 recorded repetitions per unique combination of magnetic field and heating time. Afterwards, the data for each magnetic field is shifted and linearly transformed, while preserving angles, to emphasize trend changes in a transformed coordinate system. The transformation shifts the data such that the coldest point is fixed to $(0, 0)$. Next, the data is rotated to turn the hottest point to the coordinate $(x, 0)$. Finally, the whole data is scaled by a factor $1/x$ to bring the hottest point to the coordinate $(1, 0)$. This way, the temperature is increasing from $(0, 0)$ to $(1, 0)$, making any trend changes visible as deviation from the x-axis.

A plot of the transformed latent space is shown for four exemplary magnetic fields in Fig. 5.14. Note, that each datapoint in the transformed latent space corresponds to the average position of ~ 13 recorded images. In the transformed latent space, it becomes clear that the data is not described by a single linear line, but by the adjoint combination of two piecewise linear segments.

To determine the trend changing position, we fit a piecewise linear function to the data

$$\Delta f(x) = \begin{cases} a_1(x - x_0) + b & \text{if } x \leq x_0 \\ a_2(x - x_0) + b & \text{if } x > x_0, \end{cases} \quad (5.2)$$

with a_1 and a_2 being the slopes of the two segments, and b being the value of the trend changing position along the axis of Neuron II. The function is fitted to the data in the latent space spanned by Neuron II and Neuron I, so x_0 gives the position of the trend changing position along the axis of Neuron I.

Since the data is ordered by the heating time t_{heat} , the trend changing position corresponds to a critical heating time τ_C . To determine the critical heating time τ_C , we average the heating time of the three closest datapoints to the trend changing position. Afterwards, we use the results of the thermometry in Fig. 5.2 to map the critical heating time to the critical temperature of the sample.

Statistical variation

During training, we notice that while the latent space consistently resembles the latent space from Fig. 5.13, different training iterations can produce flipped or rotated latent spaces. Although these latent spaces share a similar overall shape and can be transformed to appear more alike, they vary slightly in the precise distribution of data. These variations arise from

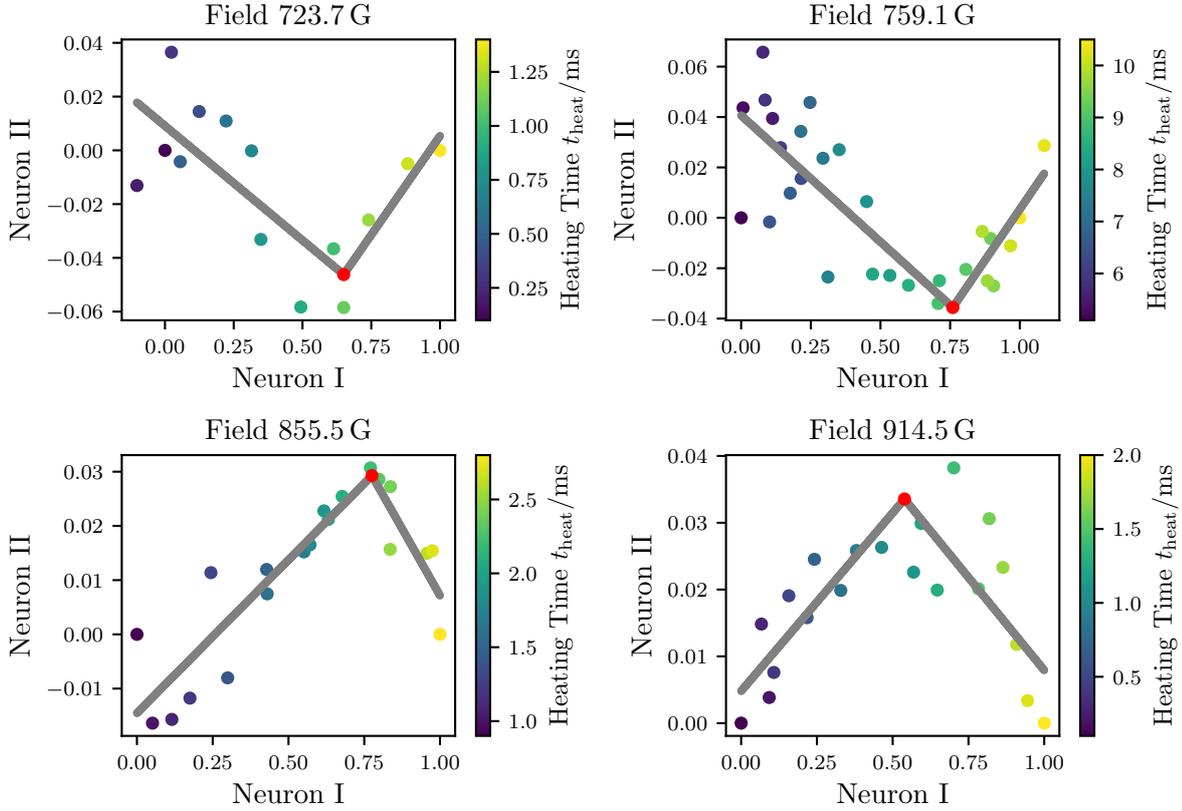


Figure 5.14: Determination of the trend changing position from the latent space. The figure depicts the transformed latent space for four exemplary magnetic fields. Each datapoint is the average position of ~ 13 images. The trend changing position is fitted by two piecewise linear segments (grey line), which are connected at the trend changing position (red dot).

the statistically variable nature of the training process, influenced by factors such as randomly initialised network weights, finite-sized and randomly shuffled training datasets, as well as stochastic gradient descent. To mitigate the effect of statistical outliers on the determined critical heating times, we train a batch of 173 models instead of relying on a single trained model.

In Fig. 5.15 we show four different transformed latent spaces for four distinctly trained models, shown for a magnetic field of 835.8 G. While the latent space generally indicates a noticeable trend changing position, some models exhibit much less pronounced trend changes. This results in a variation of the determined critical heating times τ_C .

In total we train 173 models and omit only one model due to a non-converged loss after training. Among the remaining 172 models, we discard another 33 fits out of a total of 2408 fits across all models and magnetic fields. The discarded fits have determined critical heating times at the edges of the trimmed latent space, and can not be considered trend changes. These fits mainly belong to fields close to $1/(k_F a) \approx 0.5$, where the trend change is less pronounced (refer to Fig. 5.13).

A histogram of the determined critical heating times τ_C is shown in Fig. 5.16 along with

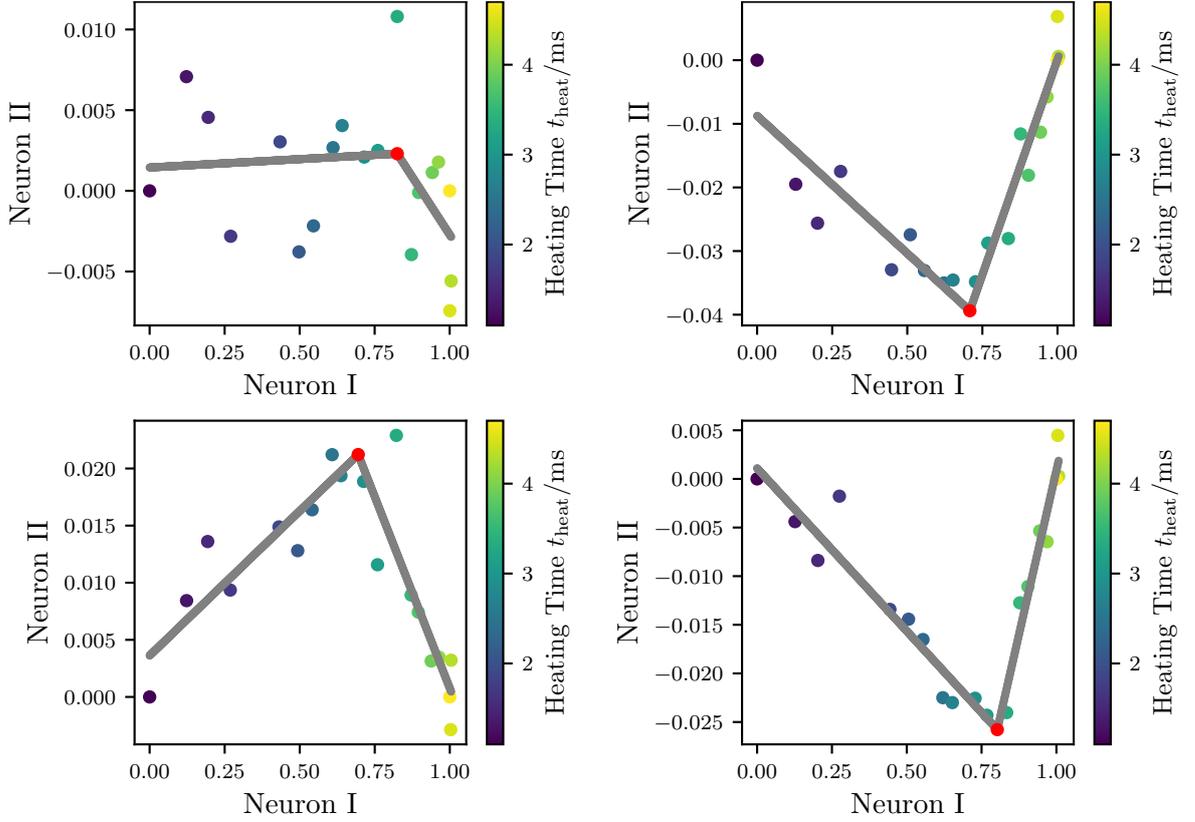


Figure 5.15: Statistical variation of the latent space after training different models. The figure depicts the transformed latent space for four distinctly trained autoencoder models, shown for the same magnetic field in each plot. Each datapoint is the average position of ~ 13 images. The trend changing position is fitted by two piecewise linear segments (grey line), which are connected at the trend changing position (red dot).

the critical heating times from the supervised analysis in Fig. 5.8. The histogram reveal a less sharp determination of the critical heating time as compared to supervised analysis in Section 5.4. However, for most fields, a specific critical heating time has a pronounced abundance. To continue the analysis, we calculate the mean and the standard error of the mean from the histograms, presenting our results for the critical heating time τ_C at each magnetic field.

An overview of the determined critical heating times is given in Fig. 5.17. While the results for most magnetic fields align closely with those from the supervised analysis, some discrepancies and larger standard errors are evident. For example, stronger deviations occur at fields closer to the BEC limit. This may be attributed to the small critical heating times of < 1 ms, shown in Fig. 5.7. At such brief heating times, trend changes at lower temperatures could be obscured by the limited amount of datapoints below the phase transition. Additional deviations occur around $1/(k_F a) \sim 0.5$, particularly at a magnetic field of 803.4 G.

The fit of the trend changing position is performed via a piecewise linear fit of two linear segments, as described by Eq. (5.2). Accordingly, the trend changing position can manifest

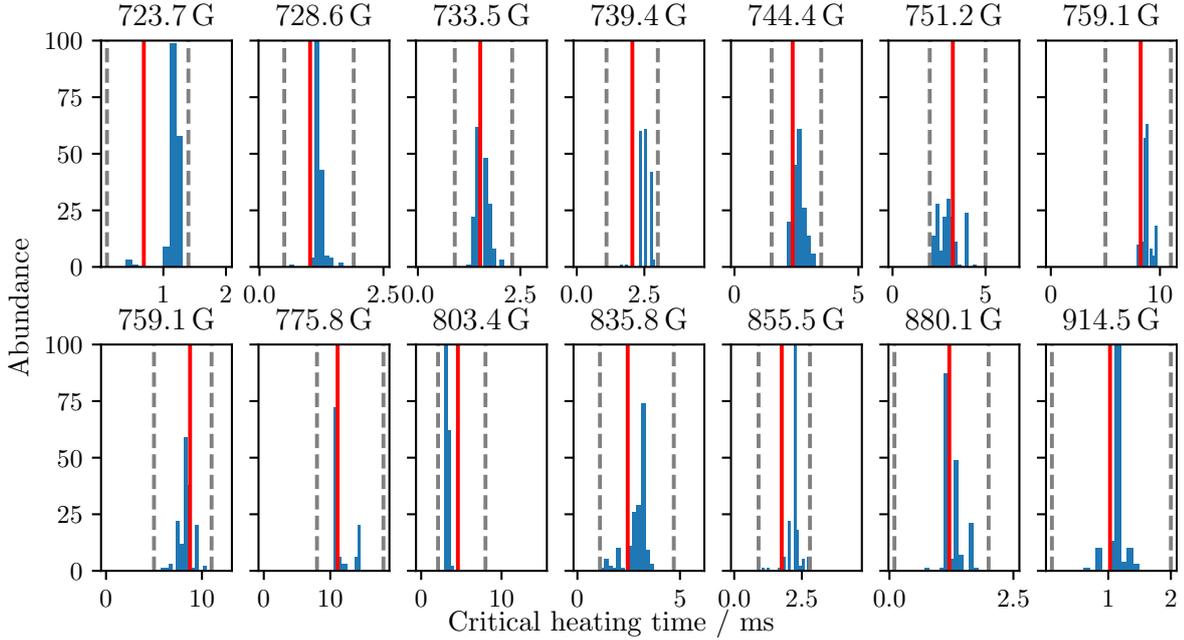


Figure 5.16: Histograms of critical heating times acquired with the autoencoder. The figure shows the abundance of each fitted critical heating time throughout the 172 analysed autoencoder models (blue bars) for each analysed magnetic field. The data range for fitting has been limited to the grey dashed lines. For comparison, the results from the supervised analysis in Fig. 5.8 are shown (red vertical line).

at positive or negative values of Neuron II. At the problematic fields, the orientation of the trend changing position changes sign, which is connected to a weakly pronounced trend change. Furthermore, these fields exhibit the smallest variation in temperature and Fermi temperature with respect to the heating time, as shown in Fig. 5.2. All of these effects can obscure the detection of trend changes.

Next, the critical heating time is used to determine the critical temperature of the superfluid phase boundary.

Trimming the data range

As shortly motivated in the previous section Section 5.5.2, the latent space is trimmed before it is transformed to the latent space shown in Fig. 5.14. Trimming is motivated by the fact, that the autoencoder has not been trained to actually detect the phase transition, but instead to compress information about the phase in the latent space. Consequently, the latent space may contain information about various features that the network deems important, not just the relatively weak signals of the second-order superfluid phase transition.

To isolate trend changes associated with the second-order phase transition, other features that could obscure this signal must be removed. Looking at Fig. 5.13, lines of equal magnetic field suggest pronounced bending for very cold samples, where the samples have a large condensate fraction, as shown in Fig. 5.7. The reason for this bending is unclear, as the autoencoder only compresses the images in a low-dimensional, abstract space. However, one

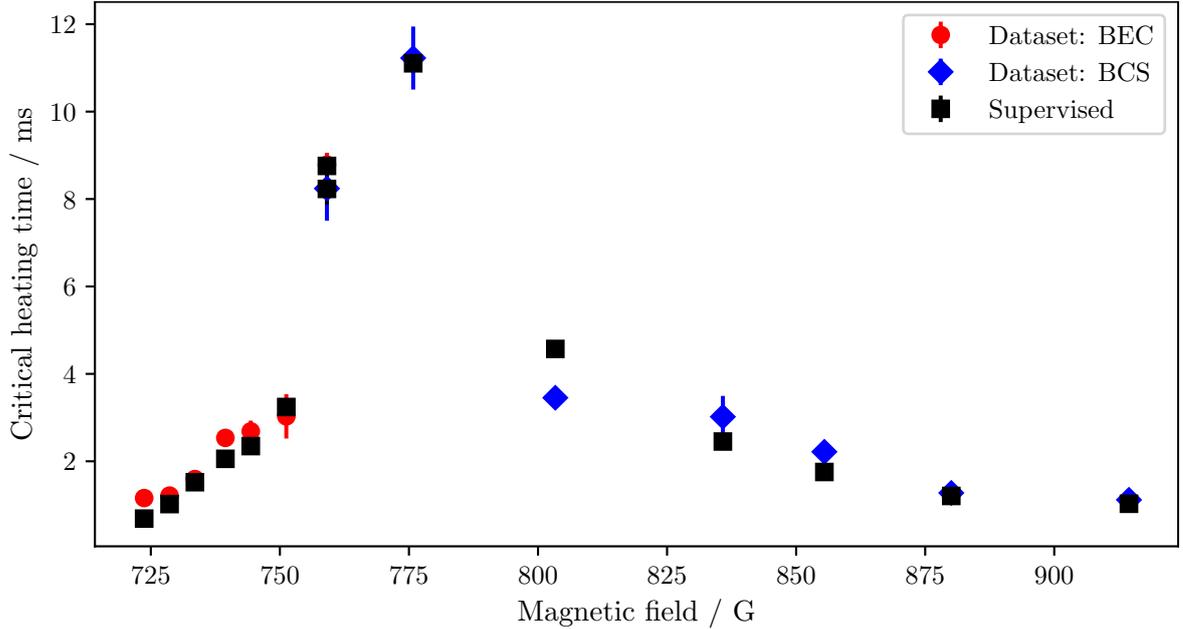


Figure 5.17: Critical heating times in the BEC-BCS crossover. The figure shows the determined critical heating times τ_C from the autoencoder analysis (red points and blue diamonds) along with the results from the supervised analysis (black squares) from Fig. 5.8. The results from the autoencoder are the mean values from Fig. 5.16 with errors denoting the standard error of the mean.

possible explanation for this bending could be the significant increase in condensate fraction at magnetic fields of 759.1 G or 775.8 G for very cold samples, as seen in Fig. 5.7.

Fig. 5.18 displays the latent space for a field of 775.8 G, where we have measured the largest critical heating time in Fig. 5.17. The figure shows the latent space in three different forms: untransformed (left), transformed without trimming (middle), and transformed and trimmed (right). The untrimmed latent space does not yield a reasonable critical heating time τ_C . Instead, the most prominent trend change occurs for very cold samples, suggesting a critical heating time around ~ 4 ms. Comparing this to the predicted CF in Fig. 5.7 suggests that this early trend change (at $CF \sim 0.1$) could be linked to deviations associated with larger condensate fractions, which also appear at < 4 ms.

By trimming the latent space around the phase transitions identified in Fig. 5.7, the fitting procedure is locked to concentrate on the trend changes connected to the superfluid phase transition.

Latent space in 3D

An open question remains regarding the architecture of the autoencoder neural network described in Table 5.2. Initially, we assume that two neurons in the bottleneck layer are sufficient for reducing the dimensionality of the input dataset. However, this could be a naive oversimplification, and further exploration is needed to justify using only two neurons in the bottleneck.

To test the influence of differently sized bottlenecks, we have trained similar networks

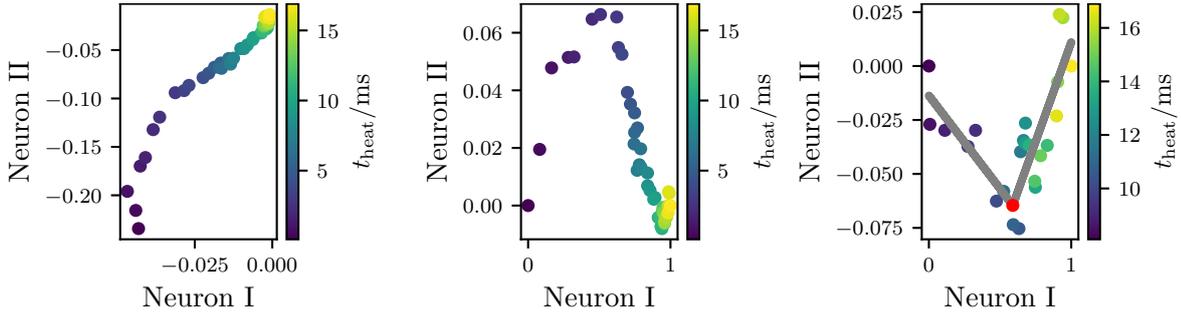


Figure 5.18: Latent space without trimming. The figure shows the latent space for an exemplary model at a field of 775.8 G. **Left:** shows the untransformed and untrimmed latent space as in Fig. 5.13. **Centre:** Transformed but untrimmed latent space. **Right:** Trimmed and transformed latent space along with a fit of the trend changing position as shown in Fig. 5.14.

but with differently sized bottleneck layers. In a first test, the bottleneck can be chosen to be just a single number or neuron. Despite this, the autoencoder converges and learns a representation capable of decoding the single value into convincing outputs. On the other extreme, we could increase the bottleneck to a size of ten neurons. Naturally, this network is also able to produce convincing outputs.

A straightforward way to quantitatively evaluate the impact of additional neurons in the bottleneck is to assess how well the output matches the input, or in other words, how small is the remaining loss after training is completed. The loss function for differently sized bottleneck layers is shown in Fig. 5.19. While a 1D bottleneck has a slightly higher loss, adding more than two neurons does not significantly reduce the final loss value. Thus, a two-dimensional bottleneck already appears sufficient to encode most of the information from the TOF images.

To further understand what features an additional dimension might capture, we have trained an autoencoder with a three-dimensional bottleneck, which otherwise exactly resembles the architecture in Table 5.2. After training, the 3D latent space is plotted in Fig. 5.20. Although the data now spans three dimensions, the latent space shows the appearance of a two-dimensional data distribution on planes in the latent space. We find that the two datasets in the neural network dataset (refer to Section 5.3.2) are actually distributed on two slightly offset but closely aligned planes. This is not unexpected, as small changes to the system can now be resolved by the third dimension. To analyse the planes, both the “BEC” and “BCS” datasets are fitted to a plane equation, respectively, using principal component analysis (“PCA”) [196].

PCA involves calculating the covariance matrix of the centred 3D data distribution. The principal axes, represented by the eigenvectors of the covariance matrix, define the directions of maximum variance. Here, the eigenvector corresponding to the smallest eigenvalue indicates the normal vector to the plane that best fits the data. The spread of datapoints away from this plane can be quantified by the ratio of the smallest eigenvalue to the norm of the eigenvalues of the other two principal axes. This results in a spread (or thickness) between 1% and 5%, depending on the training iteration.

Fig. 5.20 confirms that the datasets align well with distinct planes. After fitting the planes,

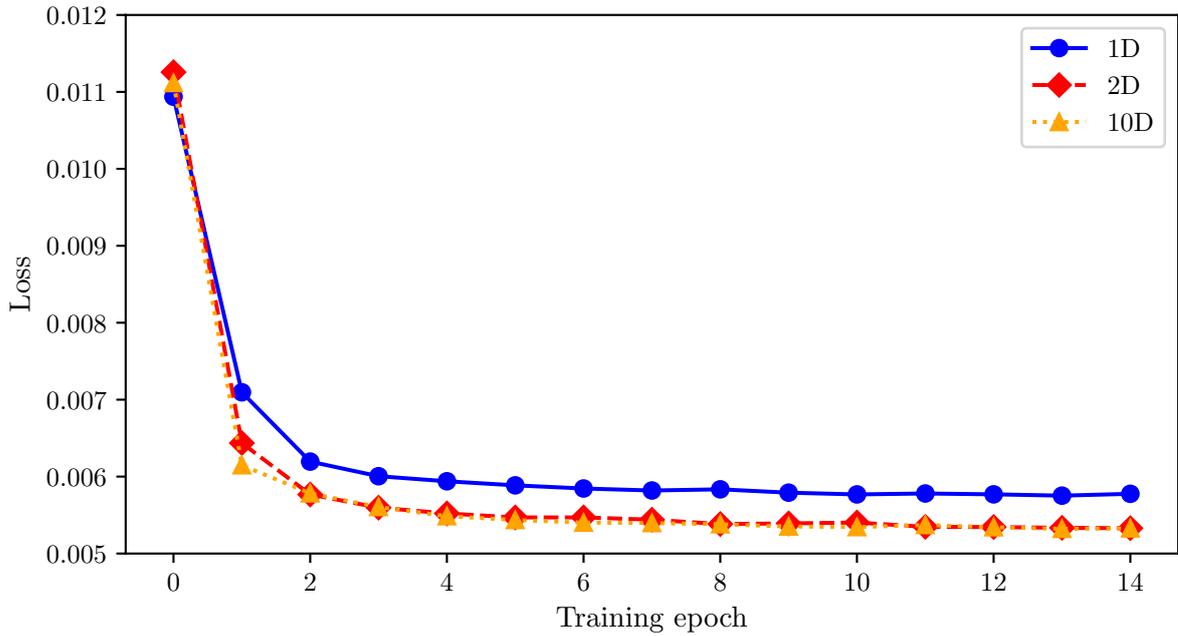


Figure 5.19: Training losses for different bottleneck sizes. The figure shows the training loss for differently sized bottleneck layers after a certain amount of training epochs. The figure has been adapted from [134].

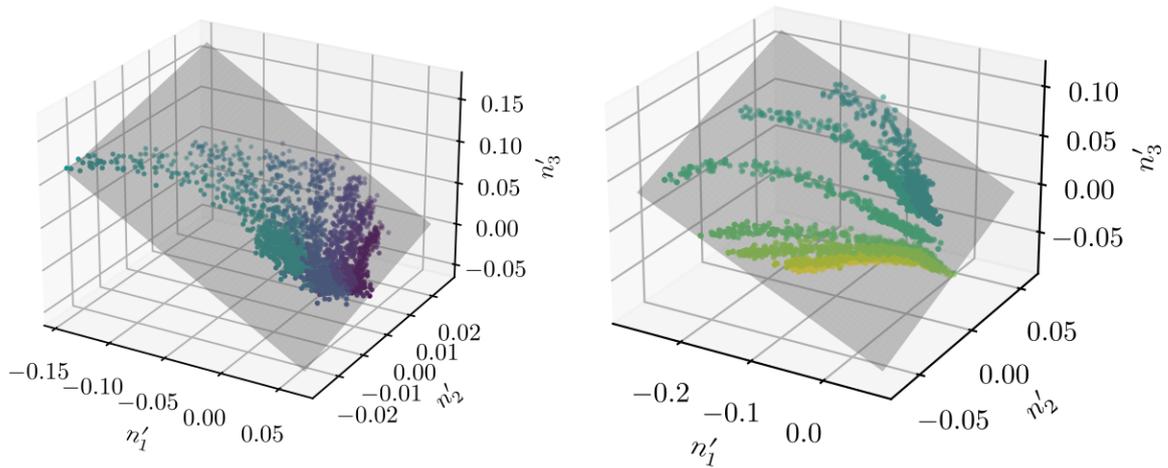
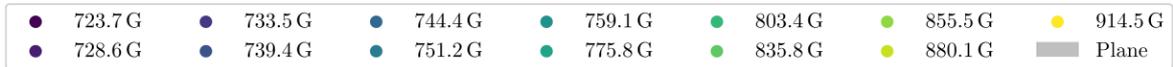


Figure 5.20: Latent space of a 3D bottleneck. The figure shows the 3D latent space from an autoencoder neural network with a 3D bottleneck layer. The used dataset is the same as in Fig. 5.13. Two distinct planes for the “BEC dataset” (left) and “BCS dataset” (right) are fitted to the 3D data distribution via PCA.

the 3D data is projected onto their respective plane, recovering a two-dimensional latent space. The exemplary projection of the “BCS dataset” is shown in Fig. 5.21. From this two-dimensional representation, the same fitting procedure as discussed in Section 5.5.2 can be used to identify the critical heating time based on trend changes as shown in Fig. 5.21.

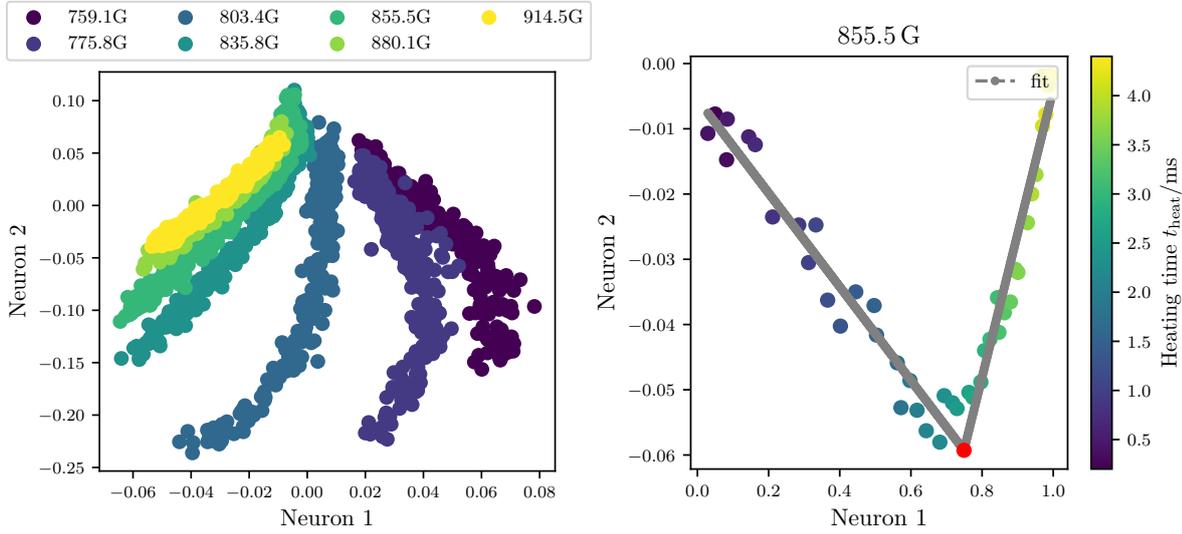


Figure 5.21: Plane projection of the 3D latent space. **Left:** The figure shows the projection of the 3D data (“BCS”) from the 3D latent space in Fig. 5.20 onto their respective plane. **Right:** Afterwards, the same fitting procedure as described in Fig. 5.14 is used to fit the critical heating time from the trend changing position.

In conclusion, a two-dimensional bottleneck suffices in capturing the trend changing features in the latent space, associated with the second-order phase transition. For higher-dimensional bottlenecks, the trend changing features can be effectively reduced to a two-dimensional plane, yielding consistent results.

5.5.3 Critical temperature

Finally, the critical heating times conducted from the trend changing positions in Fig. 5.17 can be used to calculate the critical temperature. Exactly as shown in Fig. 5.9, we calculate the critical temperature by deducing the temperature T_C and Fermi temperature T_F at the critical heating times τ_C in Fig. 5.2.

We show the determined critical temperature from the autoencoder analysis along with the results from the supervised analysis in Fig. 5.22. Our results reveal the critical temperature in the strongly interacting regime of $-0.40 < 1/(k_F a) < 0.98$. We find overall good agreement with the extended Gor’kov-Melik-Barkhudarov theory [19]. Towards the BCS limit, our results support higher temperatures, predicted by the t-Matrix approach in [45], which exceed the quantum Monte Carlo results in [47]. Around unitarity and the BEC side of the crossover, our findings align well with the quantum Monte Carlo results in [128].

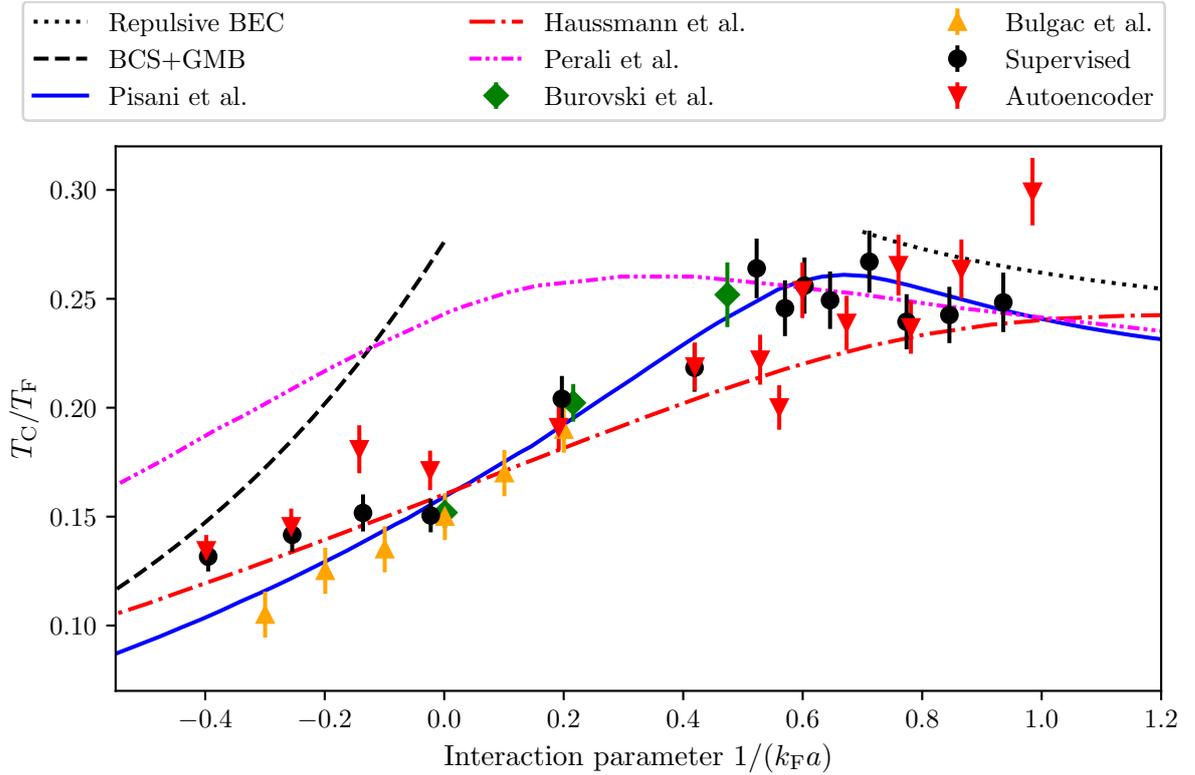


Figure 5.22: Critical Temperature in the BEC-BCS crossover via the autoencoder neural network. The figure presents the results of the superfluid critical temperature, which were determined via the autoencoder neural network (downright triangles) [68], shown alongside with our results from the (supervised) deep convolutional neural network (circles) [67]. The errorbars comprise the standard error as well as a dominating 5% uncertainty due to anharmonicities of the trapping potential (refer to Section 4.2.1). Several theoretical calculations of the critical temperature T_C/T_F are shown alongside the results. In the limits, the analytical results for the weakly interacting BEC (Eq. (2.13)) (dotted line) and the exponentially decaying BCS result with GMB correction (Eq. (2.25)) (dashed line) are shown. Across unitarity, an extended GMB correction [19] (solid line) and a self-consistent t-matrix approach [45] (dashed-dotted line) are illustrated. Additionally, two quantum Monte Carlo results (circles [128] and upright triangles [47]) are presented.

Conclusion – supervised vs. unsupervised

Overall, our findings closely match several theoretical predictions [19, 45, 128] and the results from our supervised analysis in Fig. 5.9 [133]. Therefore, we conclude that the trend changing position indeed represents a feature of the second-order phase transition, encoded by the autoencoder neural network. However, the autoencoder analysis generally displays higher uncertainty, reflecting the washed out analysis due the broader, statistical distribution of trend changing position, observed in the histograms in Fig. 5.16.

In conclusion, we successfully detect the superfluid phase transition in a strongly interacting Fermi gas from a model-free analysis, relying solely on the momentum distribution in TOF images. This marks the first detection of the superfluid phase transition across a wide range of interaction parameters in the strongly interacting regime of the BEC-BCS crossover,

achieved without the use of the rapid ramp (RR) technique. While providing an alternative perspective on the determination of the phase boundary, this approach also validates our previous methodology using the RR, and also gives confirmation for the first experimental observation of an increasing critical temperature from the BCS limit, beyond the unitarity point toward the BEC limit.

5.5.4 DeepLift in an autoencoder

In Section 5.4.4, we have employed the DeepLift library [200] to analyse the supervised deep convolutional neural network and to determine importance scores for the activation map of the input layer, that are the TOF images. The DeepLift library has allowed us to assess the physical relevance of momentum bins in TOF images to the predictions of condensate fractions CF , shown in Fig. 5.10. This analysis provides insights into how the neural network prioritises regions of the cloud that are likely influenced by pairing statistics, as illustrated in Fig. 2.6.

Here, we aim to explore whether a similar approach can enhance our understanding of the autoencoder neural network. Unlike the supervised neural network from Section 5.4.4, the autoencoder lacks the straightforward choice of an output neuron for determining the importance scores of the input neurons. In the bottleneck, two such neurons can serve this purpose. To proceed, we split the trained network into the encoder and decoder, as done in Section 5.5.2, using the encoder with the bottleneck as its output to apply the DeepLift library.

For the DeepLift analysis, we must select an output neuron. Without a clear indication of which output neuron is more significant, we conduct the analysis on both bottleneck neurons. The resulting importance scores as a function of the momentum bins k/k_F is shown in Fig. 5.23 for the coldest images, without any heating. Note, that the calculation of the momentum bins and importance scores is done exactly as introduced in Section 5.4.4.

While the interpretation of these importance scores is less straightforward than in Section 5.4.4, we aim to understand how the autoencoder differentiates superfluids from normal phase momentum distributions. In the latent space, the autoencoder cluster the coldest (condensed) samples for different interaction parameters (magnetic fields) across a wider region of the latent space, while hot (thermal) gases occupy a more concentrated and smaller area (refer to Fig. 5.13). The key question is: what pixels or regions in the TOF images most influence the autoencoder to assign an image to the region of cold (condensed) gases or the region of hot (thermal) gases?

Indeed, for the coldest samples, the network prioritises lower momenta for samples on the BEC side of the crossover, whereas for samples further on the BCS side, higher momenta gain importance, as low momenta decrease in relevance. Although the results are not as distinct as in the supervised neural network analysis in Fig. 5.10, these findings further suggest that the autoencoder learns essential (physics-based) concepts, such as temperature and interaction strength (magnetic fields), simply by extracting features from the TOF momentum distributions. Here, the DeepLift analysis suggests that the network leverages physical signatures of pairing in order to learn these concepts.

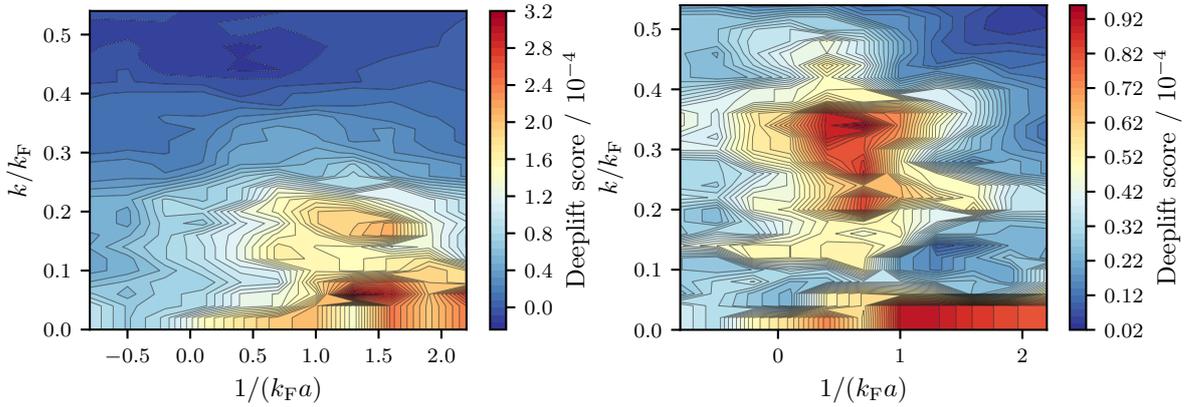


Figure 5.23: DeepLift analysis of the autoencoder neural network. The importance score of pixels in the input image is determined via the *DeepLIFT* library [200] for both bottleneck neurons (**left** and **right** image). To this end, only the coldest samples without any heating time from the training dataset are analysed. Pixels are correlated with their momentum after 5 ms of TOF, omitting interactions. Both axes are normalised with the Fermi momentum k_F , corresponding to a harmonic trap with N atoms.

5.6 Conclusion

In this chapter, we have developed two distinct neural network architectures to detect the superfluid onset in a strongly interacting Fermi gas across the BEC-BCS crossover. Both architectures infer signatures of superfluidity directly from the momentum distribution in corresponding time-of-flight (TOF) images – usually too faint for conventional fitting methods – enabling us to determine the superfluid phase diagram without relying on the conventional rapid ramp (RR) technique.

The first architecture employs a deep convolutional neural network trained in a supervised manner on labels of the condensate fraction CF , determined from corresponding measurements using the RR. During training, data near the phase transition are masked, preventing the network from overfitting to labels from the RR that suffer from low signal-to-noise ratio at a diminishing CF . After training, the network can predict the condensate fraction directly from TOF images without the RR, allowing us to determine the critical point of vanishing CF .

By measuring the in-situ 3D density distribution (refer to Section 4.2.1) at the critical point, we deduce the critical temperature T_C/T_F (refer to Section 4.2), which enables us to map the superfluid phase diagram in the BEC-BCS crossover with unprecedented accuracy. Our results experimentally show, for the first time, an increase in T_C/T_F from the BCS side beyond unitarity, with a maximum on the BEC side of the crossover. In the strongly interacting regime, our results show a striking resemblance to the extended GMB theory in [19].

Remarkably, we have found evidence that the neural network has developed a physics-based understanding of pairing using the DeepLift library [200]. An analysis with this library suggests that the network considers pixels at the centre of the cloud more important for the prediction of the CF in the BEC limit, with the relevant region shifting outwards as

one approaches the BCS side of the crossover, in line with the mechanism of pairing in the BEC-BCS crossover (refer to Fig. 2.6).

The second network architecture employs a deep convolutional autoencoder neural network – an unsupervised machine learning technique. Unlike the first network, the autoencoder does not require labelled data and instead infers information about the superfluid phase transition from structure in the available data. To this end, the autoencoder is trained to encode and decode a TOF dataset into and from a low-dimensional latent space, aiming to approximate the unity function. After successful training, the autoencoder maps different input images to distinct regions of the latent space, effectively compressing significant features of the input data into coordinates within this space.

Interestingly, the autoencoder organises data by temperature and interaction, despite only being provided with TOF images. Analysing the structure of data in the latent space reveals a change in trend around the expected superfluid critical temperature.

Deducing the temperature at these trend changing positions provides an alternative approach to constructing a phase diagram, entirely independent of the RR technique. Despite higher uncertainties, we have found very good agreement with the superfluid phase diagram obtained from our first network architecture and with the extended GMB theory in [19]. Thus, the phase diagram from the autoencoder serves as validation of our methodology, further supporting the observed increase of T_C/T_F towards the BEC side of the crossover.

Observing the Higgs mode in a strongly interacting Fermi gas

This chapter presents the results of our study on non-equilibrium effects in a perturbed, strongly interacting Fermi gas. By applying controlled perturbations in a strongly interacting Fermi gas, different modes can be excited. In the low-frequency regime, these modes are primarily driven by collective trap dynamics, such as dipole and quadrupole modes, with frequencies on the order of the trap frequency [8]. At higher frequencies, more intriguing collective modes emerge, such as the Goldstone and Higgs modes (introduced in Section 2.3.4), alongside with well-understood single-particle excitations at similar frequencies. The Higgs mode, in particular, remains poorly understood in the strongly interacting regime, largely due to the challenges involved in exciting this mode. In this thesis, we employ two distinct techniques to excite and observe the Higgs mode.

To achieve this, we examine the response of a strongly interacting Fermi gas under these two excitation methods. Both methods are implemented by a rapid change in a superimposed magnetic field, applied through the custom-built small magnetic field coil, introduced in Section 4.5.1. Using this approach, we aim to drive the Fermi gas out of equilibrium, effectively inducing the Higgs mode [73]. In general terms, a sufficiently rapid change in the magnetic field shifts the equilibrium parameters of the system faster than the gas can respond.

In the framework of a superfluid Fermi gas, the system is represented by the Mexican hat potential introduced in Section 2.3.4. After initial preparation, the system resides in an equilibrium state, characterised by the equilibrium order parameter Δ_{eq} , which corresponds to the minimum free energy of the system. Excitations of the order parameter Δ are induced by driving the system out of equilibrium, achievable by modifying the underlying Mexican hat potential. To this end, the magnetic field B is altered using the small magnetic field coil, which directly influences the scattering length a via the Feshbach resonance (refer to Section 2.2.2). This adjustment modifies the interaction parameter $1/(k_{\text{F}}a)$, effectively manipulating both the Mexican hat potential and the equilibrium order parameter (refer to Fig. 2.5). Broadly, this approach allows for two distinct excitation methods [79]: quenching or modulating the magnetic field.

The first method, illustrated on the left side of Fig. 6.1, involves a sudden quench of the magnetic field, driving the system out of equilibrium. In an early theoretical work [73], the

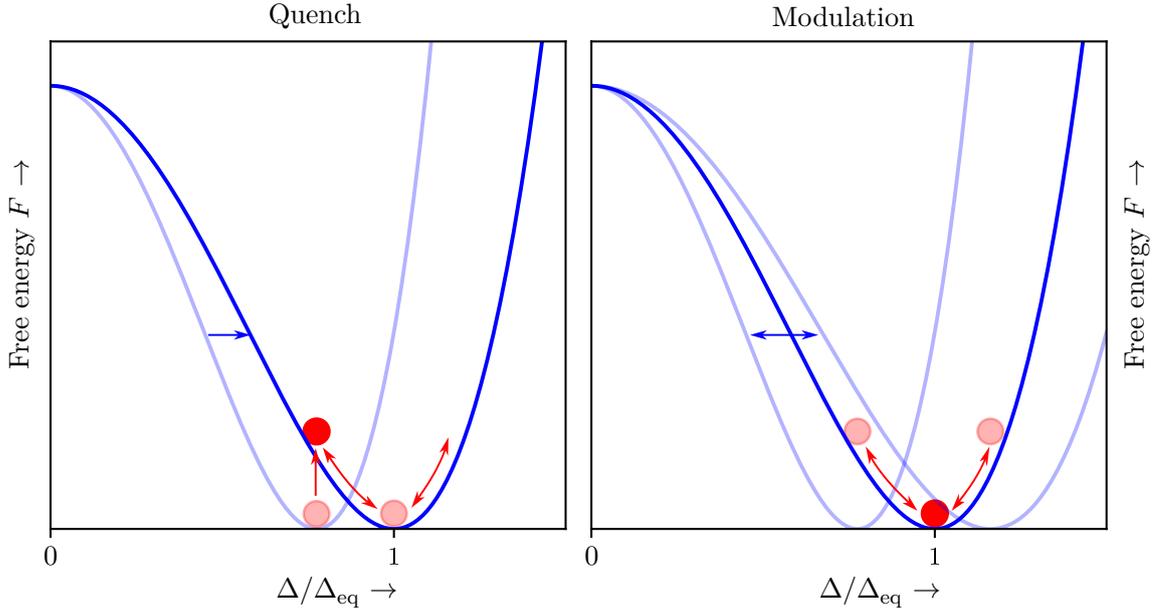


Figure 6.1: Employed excitation schemes to probe the Higgs mode in a superfluid. The sketch illustrates the two excitation schemes used in this chapter to excite a fermionic superfluid. Both panels depict the Mexican hat potential in terms of the free energy F as a function of the order parameter $\Delta/\Delta_{\text{eq}}$ for an arbitrary phase, similar to Fig. 2.8. Before introducing perturbations, the system (red dot) is described by an order parameter Δ , which is at the equilibrium position Δ'_{eq} of the Mexican hat potential, i.e. the potential minimum. Perturbations alter the Mexican hat potential and its equilibrium state (blue arrows). **Left:** Following a sudden shift in the Mexican hat potential (blue arrow), the new equilibrium value Δ_{eq} changes, but the order parameter of the system does not immediately adjust, creating a deviation $\Delta \neq \Delta_{\text{eq}}$. This deviation, shown by the lifted free energy of the system (red arrow) to an out-of-equilibrium state, induces oscillations (double headed red arrows) in the order parameter – the Higgs mode. **Right:** In a parametric excitation scheme, the Higgs mode is probed by modulating the equilibrium order parameter Δ_{eq} (blue arrow), forcing oscillation in the system (red arrow) and enabling a spectroscopic measurement of the Higgs mode.

system's response to a very small quench has been shown to be oscillations of the order parameter at twice the frequency $2\Delta/\hbar$ of the order parameter. In the same work, it has been found that these oscillations are damped with a power-law decay of $t^{-1/2}$ in the BCS limit. These oscillations have been identified later with the Higgs mode [70, 213, 214]. Exciting this mode, however, requires a quench of the system on a timescale shorter than system's intrinsic Fermi time $t_{\text{F}} = \hbar/E_{\text{F}}$, which is approximately 20 μs in our experiment. In this work, the required rapid timescale is achieved by quenching the magnetic field with the custom-built small magnetic field coil, introduced in Section 4.5.1.

The observed quench dynamics are reported in Section 6.2. Following the quench, we observe oscillations in the order parameter, manifesting as oscillations in the condensate fraction CF . From these, we derive the oscillation frequency and quality factor as a measure of the damping rate. This research has been conducted within the course of this thesis and comprises the main results of my colleague Moritz Breyer's PhD thesis [135], and hence

the experiments and analysis were mainly carried out by him. A preprint on this work is available in [101], and currently under review at SciPost Physics. My contributions include data acquisition, discussions and interpretation of the results.

The second method, illustrated on the right side of Fig. 6.1, involves parametric excitations of the order parameter by modulating the magnetic field with the custom-built small magnetic field coil. If the frequency of the modulation ν_{mod} approaches the resonance $\omega_{\text{H}} = 2\Delta/\hbar$, the Higgs mode allows for more energy to be pumped into the system. This results in a promoted decay of the condensate, which is directly observable in the condensate fraction CF after a rapid ramp, enabling a spectroscopic probe of the Higgs mode.

The results of the modulation experiments are presented Section 6.3 and are now published in [100]. This work began shortly before the works on this thesis, and was primarily conducted by my colleague Andreas Kell [134]. My contributions include data acquisition, discussions and interpretation of the results.

Outline of this chapter

In this chapter we will give an overview of the methodology and results of both methods. To put our observed results into context of the expected dynamics in the given system, the chapter is outlined as follows:

- Section 6.1: In the first section, we will present a theoretical local-density approximation (LDA) framework for the expected Higgs mode oscillations in an inhomogeneous trapping potential. In LDA, the Higgs mode is composed of many effective homogeneous oscillators, superimposing to make up the total oscillation. This presents an comparison to the homogeneous case, in which the expected Higgs mode frequency can easily be calculated. We will use this framework to judge the influence of the inhomogeneity on our observed results.
- Section 6.2: Here, we report on our results of the first proposed excitation method – quenching the superfluid. We start by introducing the experimental implementation, discuss the observed dynamics and compare our results with the developed LDA framework.
- Section 6.3: In the next section, we present our results of the second proposed excitation method – parametric excitations of the superfluid. Again, we start by introducing the experimental implementation, discuss the observed promoted decay of the condensate and compare our results within the developed LDA framework alongside a recent experimental study [75].
- Section 6.4: In the last section, we compare and discuss the results between the two employed excitation techniques – quenches and parametric excitations.

More details are given in the corresponding theses [134, 135] and publications [100, 101].

6.1 Higgs mode in an inhomogeneous trap

In most studies – unlike our experiment – the Higgs mode is examined in homogeneous systems, which are described by key parameters: the Fermi energy E_F and order parameter Δ . A possible setup for creating a homogeneous sample is presented in Appendix A, which might be employed in future endeavours (see Section 8.2.1). Parallel to this work, the Higgs mode has been investigated in an inhomogeneous sample after a quench using Bragg-spectroscopy [75]. In that approach, the focused Bragg beams probe a small, quasi-homogeneous region of the sample, effectively reducing the observation to a homogeneous Higgs mode in the centre of the trap.

In our experiment, however, we probe the entire inhomogeneous sample, a consequence of our excitation method. As proposed in [215], the inhomogeneity of the trapping potential can lead to a broader “spectrum” of Higgs mode oscillations, which can be described using a theoretical local-density approximation (LDA) framework. Within LDA (refer to Section 2.4), the Higgs mode in a homogeneous system can be mapped onto an inhomogeneous system by introducing a spatially varying Fermi energy $E_F(r)$ and order parameter $\Delta(r)$ as functions of the position r .

To explore the expected Higgs mode oscillations, we begin by considering the Higgs mode in the simplest case – a homogeneous cloud. From this basis, we apply the LDA to extend known relations to an inhomogeneous trap, as motivated by [215]. This approach is expected to yield a broadened spectrum compared to the homogeneous cloud, which may account for deviations in our observed results.

6.1.1 Homogeneous trap

The Higgs mode is identified by an oscillation of the order parameter Δ at twice its frequency, $\omega_H = 2\Delta/\hbar$ [73]. More specifically, the time evolution of the order parameter following a quench of the system’s Hamiltonian is given by the time-dependent Bogoliubov-de-Gennes equation [79, 215]. These calculations reveal an oscillation of the order parameter with a frequency given by

$$\omega_H = \begin{cases} 2\Delta/\hbar & \text{if } \mu \geq 0 \\ 2\sqrt{\Delta^2 + \mu^2}/\hbar & \text{if } \mu < 0, \end{cases} \quad (6.1)$$

with the frequency depending on the chemical potential μ if $\mu < 0$.

To estimate the frequency of the Higgs mode ω_H , the order parameter Δ of the system must be known. In a homogeneous gas at zero temperature, the Higgs mode frequency can easily be calculated using the detailed results for Δ and μ across the crossover, as given in [45] and shown in Fig. 2.5. In this scenario, the order parameter of the entire cloud oscillates uniformly.

Calculating the Higgs mode frequency becomes more complex at finite temperature $T \neq 0$. Both the order parameter Δ and chemical potential μ have a temperature dependence. In the BCS limit, only the temperature dependence of the order parameter is relevant for the Higgs mode, which is known [8] and shown in Fig. 6.2. This suggests that the Higgs mode vanishes as the gas approaches the superfluid critical temperature T_C .

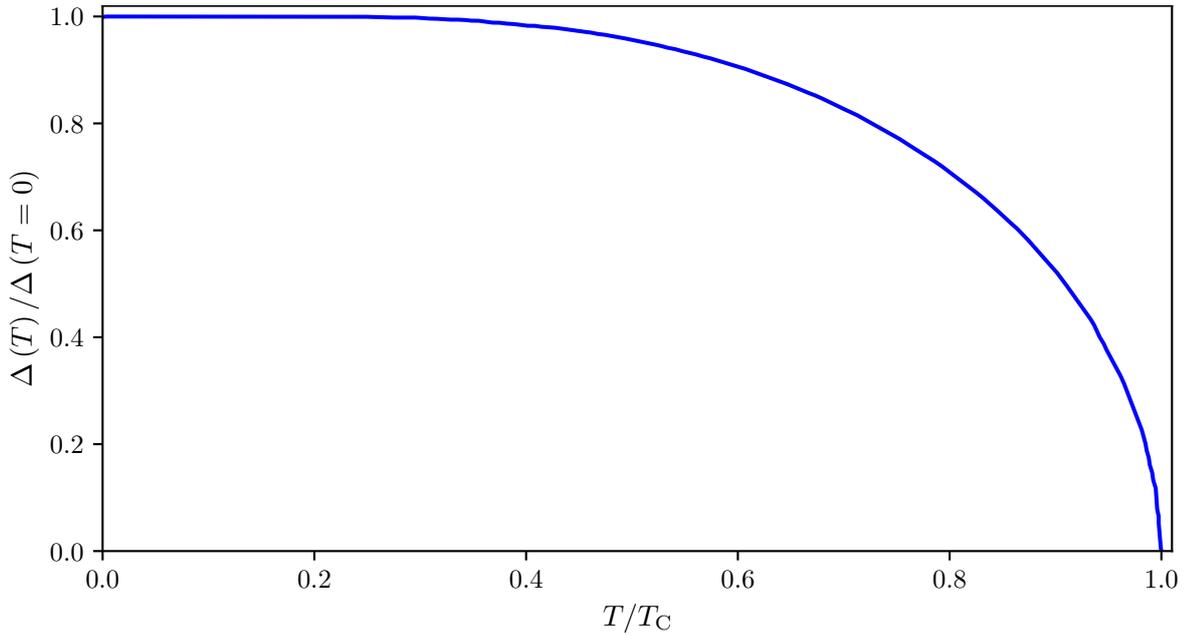


Figure 6.2: Temperature dependence of the gap parameter. The figure depicts the temperature dependence of the gap parameter $\Delta(T)$ in the BCS limit, compare to Fig. 2.5. The temperature T is given in terms of the superfluid critical temperature T_C . Data from [8].

6.1.2 Inhomogeneous trap

In an inhomogeneous trapping potential, the situation becomes notably more complex. Here, the definition of the Higgs mode in an homogeneous system is extended to the inhomogeneous system within an LDA framework, by assuming a distribution of independent Higgs mode “oscillators” throughout the trap. In this framework, each oscillator oscillates independently with its locally defined order parameter $2\Delta(r) = 2(\Delta/E_F) E_F(r)$ at the local Fermi energy $E_F(r)$, where (Δ/E_F) is given by the results in Fig. 2.5.

The integrated signal from all these oscillators constitute the full spectrum of the inhomogeneous system, resulting in a significantly broadened spectrum compared to the homogeneous gas. This broadening can be interpreted as a dephasing of the distinct oscillators, leading to a faster decay than the expected power-law decay of $t^{-1/2}$ in the BCS limit [73].

This concept, inspired by [215], has been implemented by Moritz Breyer [135] to model the Higgs mode spectrum for the recorded inhomogeneous density distributions in our experiment. This section provides an overview of the calculated model; further details are available in [135].

The spectrum $\mathcal{S}(\omega)$ of the inhomogeneous sample is given by an integral over all possible Higgs mode frequencies ω_H

$$\mathcal{S}(\omega) = \int d\omega_H g(\omega_H) \gamma(\omega, \omega_H), \quad (6.2)$$

where $g(\omega_H)$ is the density of oscillators and $\gamma(\omega, \omega_H)$ represents the linewidth of each individual oscillator. Both terms are discussed in the following parts of this section.

Density of oscillators

The density of oscillators $g(\omega_{\text{H}}) = \frac{\partial \bar{N}(\omega_{\text{H}})}{\partial \omega_{\text{H}}}$ is defined by the number of oscillators $\bar{N} = \int_0^{\omega_{\text{H}}} d\omega'_{\text{H}} N'_{\omega'_{\text{H}}}$ up to a given Higgs mode frequency ω_{H} . Here, $N'_{\omega'_{\text{H}}}$ denotes the number of oscillators in an infinitesimal frequency range around the Higgs mode frequency ω'_{H} . This is a consequence of the inhomogeneous density distribution $n(r)$, where each value of the density has a distinct Higgs mode frequency $\omega_{\text{H}}(r)$, defined by the local Fermi energy $E_{\text{F}}(r) \propto n(r)^{2/3}$ within LDA (refer to Section 2.4). Given the spatial variation of the local Higgs mode frequency $\omega_{\text{H}}(r)$, the density of oscillators can be expressed as

$$\begin{aligned} g(\omega_{\text{H}}) &= \frac{\partial \bar{N}(\omega_{\text{H}})}{\partial \omega_{\text{H}}} \\ \Leftrightarrow g(\omega_{\text{H}}) &= \frac{\partial \bar{N}(\omega_{\text{H}})}{\partial \omega_{\text{H}}} = \frac{\partial \bar{N}(r)}{\partial r} \Big|_{r=r(\omega_{\text{H}})} \frac{\partial r}{\partial \omega_{\text{H}}}. \end{aligned} \quad (6.3)$$

In the last step, the derivative has been separated into a spatial derivative of the number of oscillators $\frac{\partial \bar{N}(r)}{\partial r}$, as well as an derivative of the inverse function of the locally changing Higgs mode frequency $\frac{\partial r}{\partial \omega_{\text{H}}}$. Both of these terms can be derived from the measured density profiles $n(r)$ as $\bar{N}(r) = \int_0^r 4\pi r'^2 n_0(r')^\dagger$ and $\omega_{\text{H}}(r) = 2\Delta(r)/\hbar = 2(\Delta/E_{\text{F}}) E_{\text{F}}(r)/\hbar$, with the order parameter (Δ/E_{F}) throughout the crossover being shown in Fig. 2.5 and calculated by [45]. Note that this integral considers a density of atoms participating in the Higgs mode $n_0(r) \propto n(r)$, which is different from the density of atoms $n(r)$. This distinction will be explored further in the next part of this section. The first term of $g(\omega_{\text{H}})$ is evaluated at the position $r = r(\omega_{\text{H}})$, corresponding to the position in the trap, where the local Higgs mode frequency is given by ω_{H} .

Temperature dependence

Thus far, the density of oscillators $g(\omega_{\text{H}})$ has been considered for a zero-temperature sample, utilising the established results for the order parameter at zero temperature across the crossover, as shown in Fig. 2.5. In the experiment, however, we induce oscillations in a sample at a finite temperature, well below the superfluid critical temperature, with $T/T_{\text{C}} \approx 0.1$. As suggested by Fig. 6.2, this finite temperature introduces variations in the observed Higgs mode frequency.

We account for finite temperature in the LDA model by incorporating the known temperature dependence of the gap parameter in the BCS limit, as shown in Fig. 6.2. For this purpose, the entire cloud is assumed to be at a uniform thermalised temperature T , which is inferred from the interaction parameter $1/(k_{\text{F}}a)$ where the condensate fraction CF vanishes [20, 135]. Due to the spatially varying Fermi energy $E_{\text{F}}(r)$, this results in a locally varying reduced temperature $(T/T_{\text{F}})(r)$. At a certain position in the cloud r_{C} , the locally varying reduced temperature matches the critical temperature $T_{\text{C}}/T_{\text{F}}$ exactly. Consequently, only atoms obeying $r < r_{\text{C}}$ participate in the Higgs mode within LDA.

[†]This integral is simplified. The conducted integrals consider the surface area of equidensity shells of the elliptical density distribution

This limits the atoms participating in the Higgs mode $n_0(r) \propto n(r)$, by spatially confining them to $r < r_C$. However, this does not fully capture the whole picture of the temperature dependence. For a more complete description, the atoms involved in the Higgs mode must be weighted by a factor considering only the fraction of atoms that are part of the condensate. This matches the description of the condensate fraction CF . From an experimental point of view, the condensate fraction is treated as a global parameter, given by the integrated number of atoms in the bimodal peak (refer to Section 4.3), and thus does not effect the shape of the final Higgs spectrum $\mathcal{S}(\omega)$. In the LDA framework, however, the condensate fraction CF is considered to be spatially varying. To implement this spatial variation, the density n_0 is scaled by the squared “wave-function” of the condensate, specifically the squared order parameter $|\Delta|^2$.

Finally, the density of oscillators $g(\omega_H)$ is calculated by scaling the density of atoms $n_0(r)$ by $|\Delta(T)|^2$, obtained from Fig. 6.2. Additionally, the local Higgs mode frequency $\omega_H(r) = 2\Delta(T, r)/\hbar$ is scaled by the same temperature dependence. The resulting locally varying terms are displayed in the top row of Fig. 6.3.

Linewidth

The final component needed to calculate the Higgs spectrum $\mathcal{S}(r)$ of the inhomogeneous cloud is the linewidth $\gamma(\omega, \omega_H)$ of each individual oscillator. Already in the original work of Volkov and Kogan [73], the Higgs mode in the BCS limit has been identified to have a fundamental power-law decay of $t^{-1/2}$. This power-law behaviour has been confirmed by solution of the time-dependent Bogoliubov-de-Gennes equation [79]. Later, the damping has been calculated to approach a more damped power-law dependence of $t^{-3/2}$ in the BEC limit [131].

Calculations based on the time-dependent Bogoliubov-de-Gennes equation in [216] suggest that the transition from a BCS to BEC damping occurs relatively deep in the BEC regime, at an interaction parameter of approximately $0.5 < 1/(k_F a) < 1.5$. In our experiments described in Section 6.2 and Section 6.3, however, interactions beyond $1/(k_F a) > 0.5$ are not much investigated due to the diminishing oscillation amplitude in the condensate fraction CF . Therefore, we adopt the simpler power-law dependence of $t^{-1/2}$ to model the Higgs spectrum.

Applying a Fourier transformation to the $t^{-1/2}$ power-law dependence allows us to deduce the linewidth of a single Higgs mode oscillator. In simplified form, the linewidth of an oscillator with frequency ω_H is given by [135]

$$\gamma(\omega, \omega_H) = \sqrt{\frac{\pi}{8\omega_H|\omega - \omega_H|}}. \quad (6.4)$$

Henceforth, this linewidth is used to calculate the final spectrum $\mathcal{S}(\omega_H)$. This linewidth is also shown as the red line in the lower panel of Fig. 6.3 for comparison to the broad full spectrum.

Spectrum

The final spectrum $\mathcal{S}(\omega)$ is calculated for a given (measured) density distribution via the integral in Eq. (6.2). Using the density of oscillators $g(\omega_{\text{H}})$ along with the linewidth of each individual oscillator $\gamma(\omega, \omega_{\text{H}})$, the total spectrum is obtained and shown in Fig. 6.3.

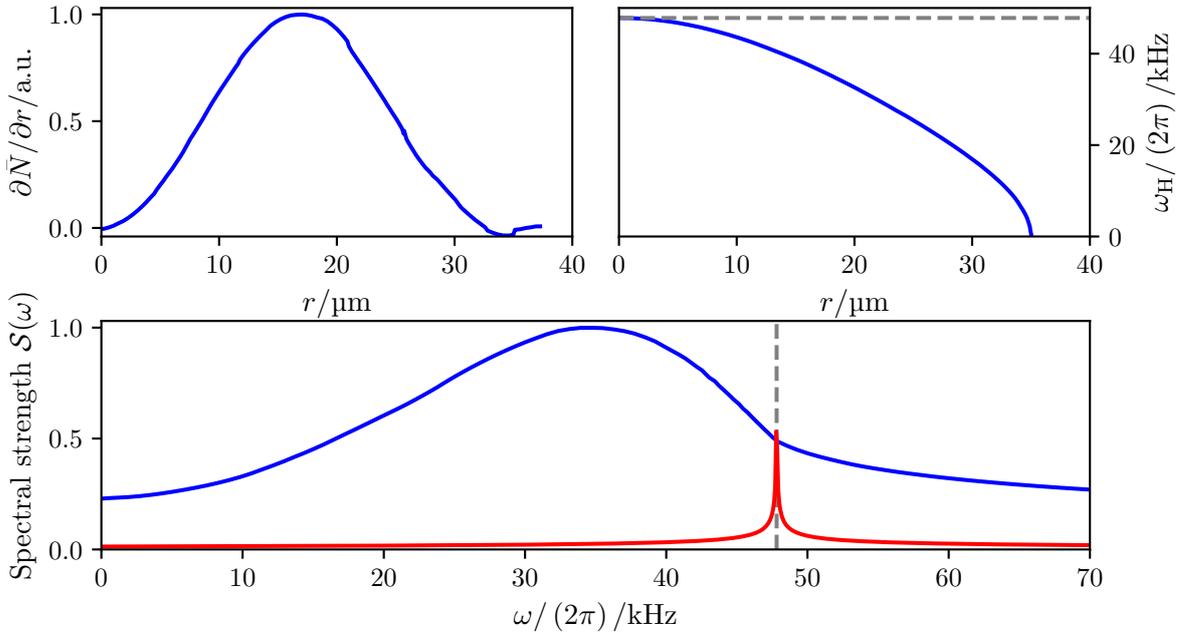


Figure 6.3: Higgs spectrum in an inhomogeneous trap within LDA. The figure depicts the exemplary calculation of a Higgs spectrum $\mathcal{S}(\omega)$ via the LDA approach given in Eq. (6.2) for a sample at unitarity $1/(k_{\text{F}}a) \approx 0.0$. As the most complete model, the shown spectrum considers finite temperature $T \neq 0$ as well as a finite linewidth of a $t^{-1/2}$ power-law dependence, given in Eq. (6.4). The local Higgs mode frequency at the centre of the sample $\omega_{\text{H}}(r=0)$ is highlighted by a grey dashed line. Additionally, the linewidth of a single oscillator at this frequency is shown by a red line and given by Eq. (6.4). **Upper left:** Spatial shell density of the spectrum, including a temperature dependence by scaling with $|\Delta(T)|^2$, as shown in Fig. 6.2. **Upper right:** Locally varying Higgs mode frequency $\omega_{\text{H}}(r) = 2\Delta(r)\hbar$, considering the same temperature dependence. **Lower:** Resulting spectrum $\mathcal{S}(r)$ (normalised) of the inhomogeneous (measured) density distribution. Figure adapted from [135].

The spectrum is shown for an exemplary magnetic field at unitarity in the lower panel as a blue line. For reference, the spectrum also includes the Higgs mode frequency corresponding to the centre of the trap $\omega_{\text{H}}(r=0)$, as a grey dashed line, providing a benchmark for the maximum expected frequency contributing to the overall spectrum $\mathcal{S}(\omega)$. The resulting spectrum illustrates that, due to the integration over all contributing local Higgs oscillators, the centre of the spectrum shifts to significantly lower frequencies, as anticipated.

Additionally, the figure displays the linewidth of a single Higgs mode oscillator at the centre Higgs mode frequency (see Eq. (6.4)), shown in red. A comparison of the linewidth of the resulting spectrum with the power-law decay reveals a much broader spectrum due to dephasing of the many partaking oscillators.

Fourier transformation of the spectrum

The resulting spectrum is then Fourier-transformed to provide an expectation for the temporal oscillations observed in Section 6.2, taking into account the trap's inhomogeneity. In the experiment, we study various quenches with different initial $1/(k_F a_i)$ and final $1/(k_F a_f)$ interaction parameters. Following these quenches, no immediate change in the density distribution is observed, and therefore, the spectrum is calculated using the measured density distribution at the initial interaction parameter $1/(k_F a_i)$ [135]. The Fourier-transformed spectra for different exemplary quenches are shown in Fig. 6.4.

For all displayed quenches, the shown temporal oscillations exhibit strong damping, with the second peak already heavily suppressed. This results from the spectrum broadening or dephasing, leading to a much more rapid decay than the $t^{-1/2}$ power-law decay in the BCS limit [215].

The least amount of discernible oscillations are seen on the BCS side of the crossover, which is a consequence of the lower Higgs mode frequency (refer to Fig. 2.5). Towards the BEC side, more oscillations become discernible, as the higher frequency allows for more oscillations to fit into the same time frame before dephasing reduces the amplitude too much. In the experiment, however, oscillations on the BEC side are expected to be significantly suppressed due to the vanishing of particle-hole symmetry on the BEC (refer to Section 2.3.4), an effect not considered in the LDA model.

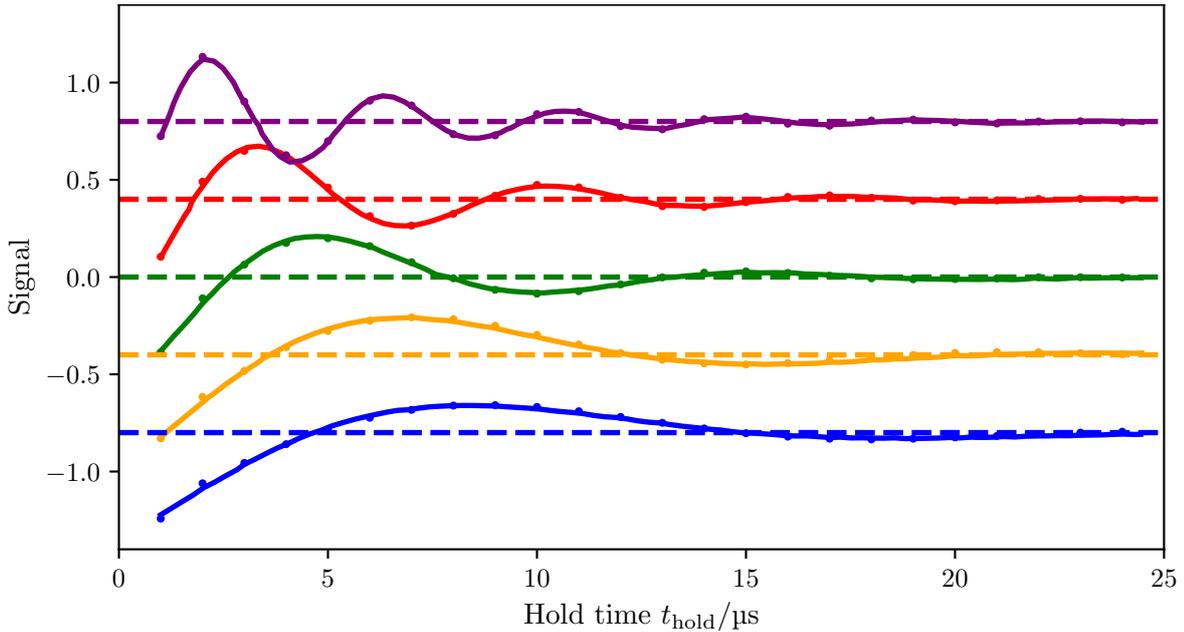


Figure 6.4: Time dynamics from the LDA Higgs spectrum. The figure presents the ambiguous signal (value) of the obtained Fourier transformations of the resulting spectra calculated in Fig. 6.3. Five different time evolutions (dots) are shown for quenches of the interaction parameter $1/(k_F a_i) \rightarrow 1/(k_F a_f)$, comprising (from top to bottom) $0.27 \rightarrow 1.08$, $0.15 \rightarrow 0.78$, $0.04 \rightarrow 0.54$, $-0.06 \rightarrow 0.4$, $-0.14 \rightarrow 0.24$. The signals for different quenches are vertically offset for clarity. Additionally, an exponentially decaying sine wave is fitted to the resulting oscillations (solid lines), which yields the expected oscillation frequency. Figure adapted from [135].

In the end, the simulated (LDA) Higgs mode in both the frequency (refer to Fig. 6.3) and time domains (refer to Fig. 6.4) provides an additional benchmark of the expected spectrum and temporal oscillation to compare with our experimental data in Section 6.2 and Section 6.3. This allows for a discussion of how inhomogeneity influences the observed Higgs mode measurements, which may explain deviations to the expected Higgs mode behaviour in a homogeneous system.

6.2 Quenching a superfluid - Time dynamics of the Higgs mode

This section presents the results of the quenches performed on the order parameter, resulting in time-resolved oscillations of the condensate fraction, a manifestation of the Higgs mode. As previously mentioned, this project has primarily been conducted by Moritz Breyer [135], while my contributions comprise data acquisition, discussions and interpretation of the results. A more detailed discussion and presentation of the results presented here can be found in Moritz Breyer's PhD thesis [135] and in the preprint of this work [101], which is currently under review.

A quench refers to a sudden change in the system's parameters. If this change occurs faster than the system's fastest intrinsic timescale, the quench can be regarded as instantaneous. This timescale is the Fermi time \hbar/E_F , which is approximately 20 μs in our experiment. Using the custom-build small magnetic field coil, described in Section 4.5.1, we can change the magnetic field faster than this timescale, resulting in an effectively instantaneous change for the system. Since the magnetic field is linked to the scattering length a (see Section 2.2.2), this change translates into a shift in the interaction parameter $1/(k_F a)$, which corresponds to a quench in the order parameter Δ (refer to Fig. 2.5).

The expected dynamics induced by this quench have previously been discussed in Section 6.1, with a more general description of the Higgs mode provided in Section 2.3.4. Here, we briefly revisit the key aspects of these dynamics to frame the experiment within the theoretical context. In the BCS limit, the Higgs mode has been identified as a damped oscillation of the order parameter at twice its frequency $2\Delta/\hbar$ and with a power-law damping of $t^{-1/2}$ [73]. More broadly, a quench from an initial order parameter Δ_i to a final order parameter Δ_f can be described using the time-dependent Bogoliubov-de-Gennes equation. After the quench, the new non-equilibrium order parameter Δ_∞ typically differs from its equilibrium value at the final interaction parameter $\Delta_\infty \neq \Delta_f = \Delta|_{1/(k_F a_f)}$ [79]. Yet, oscillations appear at twice the value of the new order parameter $2\Delta_\infty/\hbar$ [79, 216, 217], with damping described by $t^{-1/2}$ in the BCS limit and $t^{-3/2}$ in the BEC limit [131], linked by a crossover region [216].

A more detailed description about possible quench dynamics is encapsulated in the quench phase diagram shown in Fig. 6.5 [217]. The phase diagram identifies four distinct quench regimes, labelled I, II, II' and III, each exhibiting conceptually different dynamics following a quench. For very large quenches from a high initial order parameter Δ_i to a much lower final order parameter Δ_f , regime I applies, resulting in a vanishing order parameter $\Delta_\infty \rightarrow 0$, with no oscillations expected in this regime. More moderate quenches fall within regimes II and II', characterised by power-law damped oscillations with frequency $2\Delta_\infty$, after which the order parameter settles at a new equilibrium value of $\Delta(t)|_{t \rightarrow \infty} \rightarrow \Delta_\infty$. Regimes II and II' describe quenches to a final order parameter Δ_f with a positive chemical potential

$\mu_f > 0$ (II) or negative chemical potential $\mu_f < 0$ (II'). For quenches ending in the BEC regime with $\mu_f < 0$ (II'), the damping follows a stronger power-law decay of $t^{-3/2}$ [131] rather than the weaker $t^{-1/2}$ in regime II [216], consistent with the previously discussed behaviour for quenches on the BEC side. Regime III would be particularly interesting, as it features persistent oscillations of the order parameter; however, this regime is not accessible with our current methods at this time.

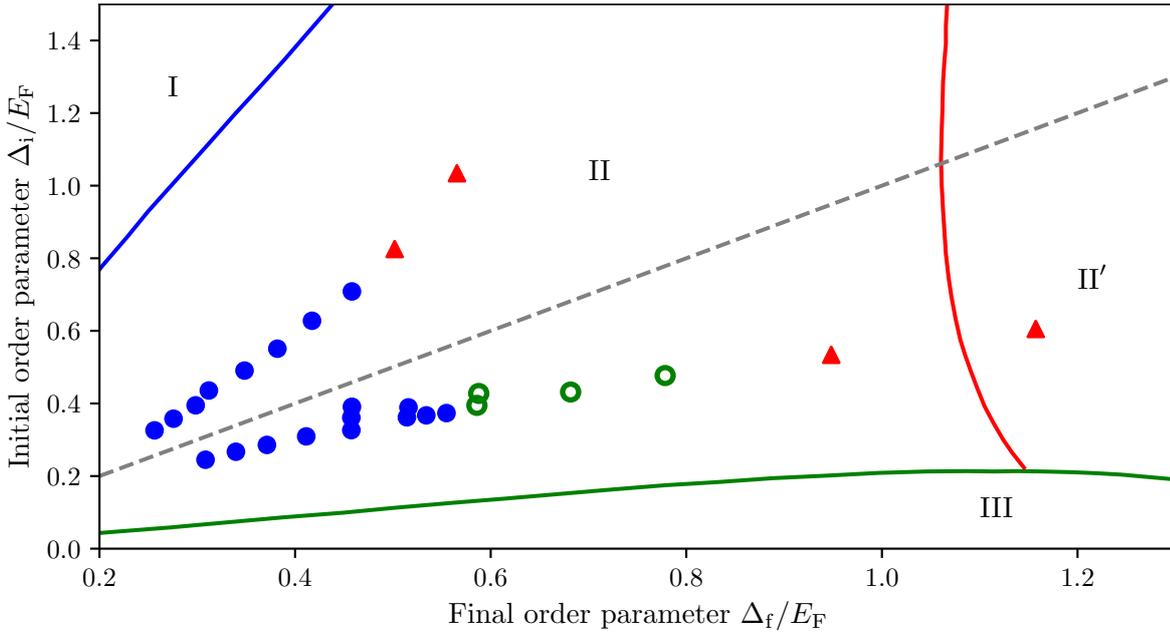


Figure 6.5: Phase diagram of quenches in a Fermi superfluid. The figure shows the quench phase diagram for quenches of the order parameter from an initial value Δ_i/E_F to a final value Δ_f/E_F . Here, $E_F \propto n(r=0)^{2/3}$ corresponds to the Fermi energy calculated from the density $n(r=0)$ at the centre of the trap. Regimes of different quench dynamics are separated by solid lines and labelled with Roman numerals I, II, II' and III. The dashed line denotes an infinitesimal quench of $\Delta_f = \Delta_i$. Quenches performed in this thesis are marked with coloured datapoints. Solid (blue) circles represent quenches with discernible decaying oscillations, while open circles (green) represent a time evolution consistent with dephasing in the LDA model (see Section 6.1). Triangles (red) denote quenches where the time evolution can not be explained by dephasing within the LDA model alone. Figure and datapoints adapted from [135]. Phase diagram boundaries from [217].

In this section we probe regimes II and II' through appropriate quenches in the magnetic field. Quenches probed in the following parts of this section are highlighted by datapoints in Fig. 6.5, with colour (and marker style) denoting the kind of observed dynamics following the quench. Discernible oscillations have only been observed for a subset of quenches with relatively small initial and final order parameters. For quenches further towards the BEC side, no discernible oscillations have been detected, which could be explained by the LDA model of the Higgs mode in an inhomogeneous trap, as introduced in Section 6.1. For quenches furthest on the BEC side, with the highest expected Higgs mode frequency, no discernible oscillations have been observed. This can not be accounted for by a faster decay due the dephasing within the LDA model and instead suggest a vanishing of the Higgs mode towards

the BEC limit, as anticipated from the loss of particle-hole symmetry (see Section 2.3.4).

The experimental protocol and resulting observations are discussed in the next parts of this section.

6.2.1 Experimental realisation

The quench of the order parameter is driven by a quench in the magnetic field. As previously discussed, the magnetic field B defines the scattering length a via the corresponding Feshbach resonance in Section 2.2.2. By varying the scattering length a , the equilibrium value of the system's order parameter Δ changes due to the shift in the interaction parameter $1/(k_F a)$. The initial and final order parameters can be calculated from the interaction parameter values in Fig. 2.5.

In the experiment, large quenches are preferred, as these they provide a higher signal-to-noise ratio in terms of changes in the condensate fraction CF . The most significant change in the system properties occurs when a small adjustment in the magnetic field B results in a substantial shift in the scattering length a . Therefore, we use the narrower $|13\rangle$ resonance of the broad Feshbach resonances, which causes a stronger variation in a with respect to B .

The rapid variation in the magnetic field is applied with the small magnetic field coil, introduced in Section 4.5.1. A schematic of the magnetic field ramps used for the quench is shown in Fig. 6.6. We prepare a sample at a magnetic field that is the superimposed sum of the Feshbach coil's magnetic field B (refer to Section 3.2.6) and the magnetic field B_{quench} from the custom-built small magnetic field coil. Typically, the quench amplitude is set to $B_{\text{quench}} = 35 \text{ G}$ [135], although it can be reduced to examine smaller quench steps. By reversing the polarity of the small magnetic field coil, the preparation field can be set to $B \pm B_{\text{quench}}$. At time $t = 0$, the quench is performed by turning off the small magnetic field coil, and depending on its polarity, the magnetic field undergoes a quench from $B \pm B_{\text{quench}} \rightarrow B$. The timescale of the quench is $\sim 2.6 \mu\text{s}$ [163] (refer to Section 4.5.1), which is much faster than the system's Fermi time of about $20 \mu\text{s}$. Thus, the quench can be considered instantaneous, and the now out-of-equilibrium system begins to evolve.

Immediately following the quench, the density distribution remains unchanged, but the order parameter deviates from its new equilibrium value, as defined by the updated interaction parameter. After a hold time t_{hold} , the dipole trap is turned off, and the system's condensate fraction CF is measured using the RR technique, introduced in Section 4.3. The condensate fraction CF has a one-to-one correspondence with the order parameter and presents an accessible proxy, serving as our observable after an evolution time of t_{hold} .

6.2.2 Observed dynamics

The experiment is conducted for both quench directions $B \pm B_{\text{quench}}$. For a positive quench $B + B_{\text{quench}} \rightarrow B$, the interaction parameter shifts to higher values $1/(k_F a_i) < 1/(k_F a_f)$, corresponding to a larger final order parameter $\Delta_i < \Delta_f$. For the opposite quench direction $B - B_{\text{quench}} \rightarrow B$, the interaction parameter shifts to lower values $1/(k_F a_i) > 1/(k_F a_f)$, corresponding to a smaller final order parameter $\Delta_i > \Delta_f$. The condensate fraction CF after various quenches $1/(k_F a_i) \rightarrow 1/(k_F a_f)$ in either direction and a varying hold time t_{hold} is shown in Fig. 6.7.

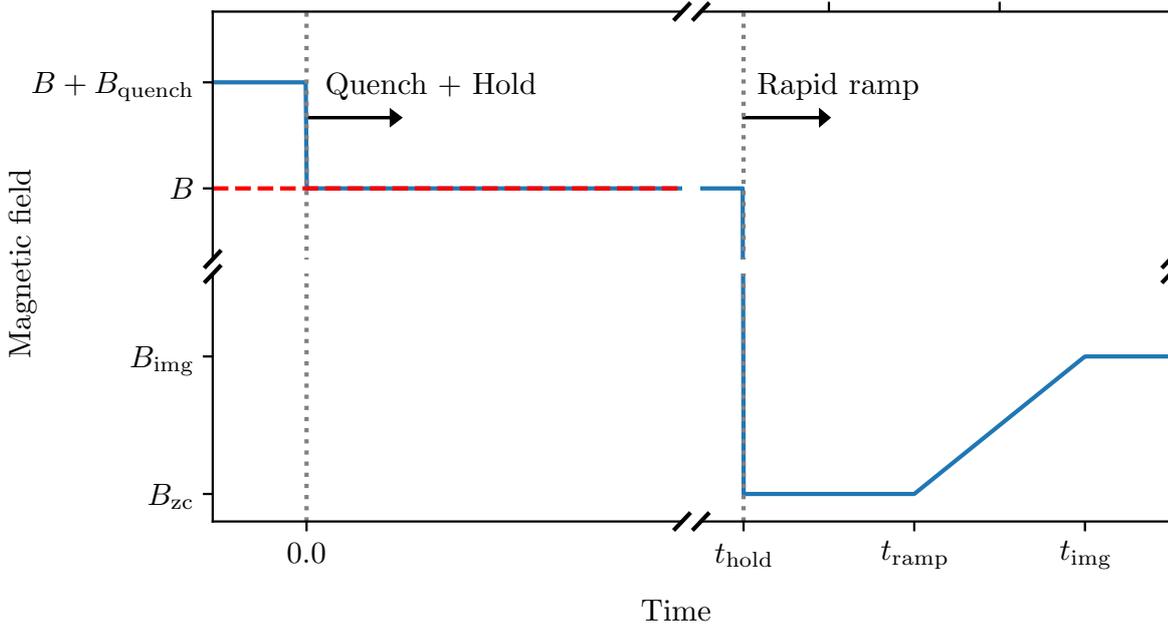


Figure 6.6: Experimental protocol to drive quenches. The figure shows the magnetic field ramps used to investigate quenches of the interaction parameter $1/(k_{\text{F}}a)$. First, the sample is prepared at a magnetic offset field of B along with a superimposed magnetic quench field B_{quench} via the small magnetic field coil, introduced in Section 4.5. This generates a preparation magnetic field of $B \pm B_{\text{quench}}$, which is shown for the positive $+B_{\text{quench}}$ direction in the plot. At $t = 0$, the small magnetic field coil is turned off, which returns the magnetic field back to B on a timescale faster than the Fermi time \hbar/E_{F} . The quench induces oscillations in the order parameter of the trapped Fermi gas. After a hold time t_{hold} , the trap is turned off, and the condensate fraction CF is measured using the RR technique, introduced in Section 4.3. The rapid ramp involves two timescales: the free expansion time $t_{\text{ramp}} - t_{\text{hold}}$ at the zero crossing of the Feshbach resonance B_{zc} , and the ramp time $t_{\text{img}} - t_{\text{ramp}}$ to the final imaging magnetic field B_{img} .

The observed quench dynamics exhibit qualitatively different time evolutions depending on the performed quench. For quenches towards the BEC side ($1/(k_{\text{F}}a_{\text{i}}) < 1/(k_{\text{F}}a_{\text{f}})$), we observe the most discernible yet heavily damped oscillations for quenches that both starts and end on the BCS side $1/(k_{\text{F}}a_{\text{i}}) < 0 < 1/(k_{\text{F}}a_{\text{f}})$. For quenches that transition from the BCS to the BEC side $1/(k_{\text{F}}a_{\text{i}}) < 0 < 1/(k_{\text{F}}a_{\text{f}})$, no clear oscillation are discernible; instead, the condensate fraction CF shows a single prominent peak at short hold times t_{hold} , followed by a subsequent decay. When both the initial and final interaction parameters are on the BEC side $0 < 1/(k_{\text{F}}a_{\text{i}}) < 1/(k_{\text{F}}a_{\text{f}})$, the initial peak vanishes, and the condensate fraction CF decays immediately.

Quenches in the opposite direction towards the BCS side $1/(k_{\text{F}}a_{\text{i}}) > 1/(k_{\text{F}}a_{\text{f}})$ display similar behaviour. When both the initial and final interaction parameters are on the BEC side $0 < 1/(k_{\text{F}}a_{\text{f}}) < 1/(k_{\text{F}}a_{\text{i}})$, only a strong decay is observed, with no discernible oscillations. For quenches entering the BCS regime $1/(k_{\text{F}}a_{\text{f}}) < 0 < 1/(k_{\text{F}}a_{\text{i}})$, however, damped oscillations become apparent, qualitatively reflecting the observations for the opposite quench direction.

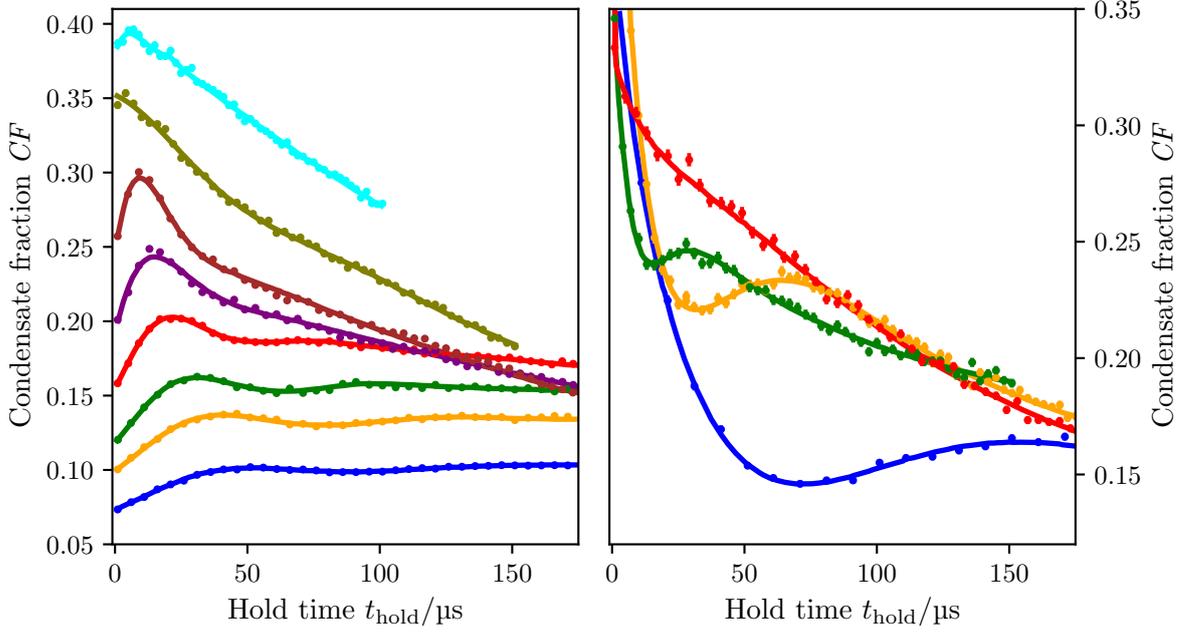


Figure 6.7: Observed quench dynamics. The figure presents the observed dynamics in the condensate fraction CF after a quench from an initial interaction parameter $1/(k_{\text{F}}a_i)$ to a final interaction parameter $1/(k_{\text{F}}a_f)$, and a subsequent hold time t_{hold} , before measuring the CF via the RR (refer to Section 4.3). **Left:** Observed dynamics for quenches with $1/(k_{\text{F}}a_i) < 1/(k_{\text{F}}a_f)$, ranging from (from top to bottom) $0.27 \rightarrow 1.08$, $0.15 \rightarrow 0.78$, $-0.06 \rightarrow 0.4$, $-0.14 \rightarrow 0.24$, $-0.22 \rightarrow 0.11$, $-0.31 \rightarrow 0.00$, $-0.37 \rightarrow -0.10$ and $-0.44 \rightarrow -0.20$. **Right:** Observed dynamics for quenches with $1/(k_{\text{F}}a_i) > 1/(k_{\text{F}}a_f)$, covering $0.90 \rightarrow 0.21$ (red), $0.44 \rightarrow 0.00$ (green), $0.18 \rightarrow -0.18$ (orange) and $-0.05 \rightarrow -0.36$ (blue). Errorbars denote the standard error from 50 repetitions per data point, which may be obscured by the symbols. Data adapted from [135].

Fit model

To quantitatively analyse the observed dynamics, we fit the data with a phenomenological model to extract the oscillation frequency and damping factor. The model function $F(t)$ incorporates a damped oscillation $f(t)$ superimposed on an empirical background function $b^{(l)}(t)$ that captures the observed decay, and reads

$$F(t) = \begin{cases} f(t) + b(t) & \text{if } 1/(k_{\text{F}}a_f) > 1/(k_{\text{F}}a_i) \\ f(t) + b'(t) & \text{if } 1/(k_{\text{F}}a_f) < 1/(k_{\text{F}}a_i). \end{cases} \quad (6.5)$$

The empirical background is chosen based on the quench direction $1/(k_{\text{F}}a_f) \lesseqgtr 1/(k_{\text{F}}a_i)$ to capture the qualitatively different observations in Fig. 6.7. For quenches towards the BEC side $1/(k_{\text{F}}a_f) > 1/(k_{\text{F}}a_i)$, the empirical background $b(t)$ includes an initial exponential approach to a new equilibrium condensate fraction, followed by a continuous decay. The background

function $b(t)$ reads

$$b(t) = \begin{cases} \eta_+ \left(1 - e^{-t/\tau_+}\right) + \eta_0 & \text{if } t \leq t_{\max} \\ e^{-(t-t_{\max})/\tau_-} \left(\eta_+ \left(1 - e^{-t/\tau_+}\right) + \eta_0\right) & \text{if } t > t_{\max}, \end{cases} \quad (6.6)$$

where the initial exponential approach from an initial condensate fraction η_0 to an increased condensate fraction $\eta + \eta_+$ occurs over a time t_{\max} with a time constant τ_+ . This mirrors the behaviour of the CF approaching a new equilibrium value, as expected in region II of the phase diagram in Fig. 6.5. After reaching $t > t_{\max}$, the condensate fraction decays exponentially with a time constant τ_- , attributed to thermal excitations gradually destroying the condensate. During the fitting process, the time t_{\max} is not constrained to lie within the measured hold times t_{hold} .

For quenches towards the BCS side $1/(k_F a_f) < 1/(k_F a_i)$, the observed time dynamics in Fig. 6.7 suggest a different empirical background function $b'(t)$, described by the sum of a fast and a slow exponential decay. This background reads

$$b'(t) = \eta_1 e^{-t/\tau_1} + \eta_2 e^{-t/\tau_2} + \eta_0, \quad (6.7)$$

where the first two terms denote the two exponential decays with distinct time constants $\tau_{1,2}$ and decay amplitudes $\eta_{1,2}$ from an initial condensate fraction η_0 .

With the background now modelled, the remaining oscillatory component can be fitted with a damped oscillation to extract the oscillation frequency ω and damping time constant τ . As identified in the BCS limit [73], oscillations of the order parameter exhibit a power-law damping of $t^{-1/2}$. However, for a more robust fitting routine, the model function instead uses an exponential decay, which reads

$$f(t) = A e^{-t/\tau} \cos(\omega t + \phi), \quad (6.8)$$

where A is the oscillation amplitude and ϕ is the initial phase of the oscillation.

Long-time dynamics

The observed time dynamics in Eq. (6.5) occur on timescales from $\sim 10 \mu\text{s}$ to $\sim 100 \mu\text{s}$, corresponding to frequencies of approximately $\sim 10 \text{ kHz}$ to $\sim 100 \text{ kHz}$. While such dynamics could potentially originate from simpler collective trap modes [8], these modes generally have lower frequencies on the order of the trap frequency $\sim 200 \text{ Hz}$, which does not affect the dynamics on the timescales considered in this experiment. In the experiment, we observe these trap modes, which can readily be distinguished from the much faster Higgs mode by their significantly slower timescales. As such, no changes in density distribution are observed over the timescales relevant to this experiment.

Consequently, no atom loss has been detected during the first 10 ms following the quench. Atom loss only becomes apparent on the timescales of the trap dynamics. Therefore, for all quenches considered, the initial and final interaction parameters are described by the same Fermi energy E_F . This has also been exploited for the calculation of the expected oscillations within the LDA model (refer to Section 6.1).

6.2.3 Extracted oscillation frequency and damping

The complete model $F(t)$ from Eq. (6.5) is fitted to the observed time dynamics in Fig. 6.7 to extract both the oscillation frequency ω and the damping time constant τ . Additionally, we also apply the damped oscillation model $f(t)$ from Eq. (6.8) to fit the simulated data from the LDA model shown in Fig. 6.4. The simulated quenches are calculated for the same $1/(k_F a_i)$ and $1/(k_F a_f)$ as the experimentally conducted quenches, incorporating the measured density profile (see Section 4.2.1) from before the quench, as well as the initial a_i and final a_f scattering lengths derived from the corresponding magnetic fields. Both the results from the experimental data and the LDA model are shown in Fig. 6.8. To normalise the axes, the observed frequencies and interaction parameters are normalised by the Fermi energy E_F derived from the density at the trap centre (see Section 2.4).

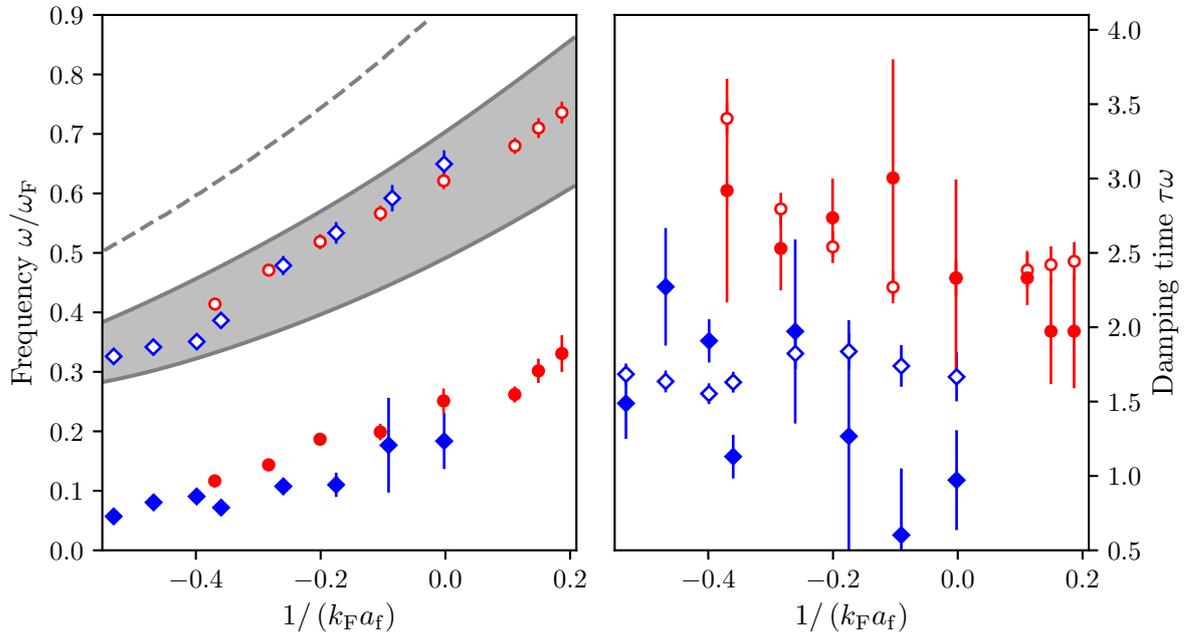


Figure 6.8: Oscillations frequencies and damping of the observed dynamics. The figure presents the extracted frequencies ω and damping times τ in terms of the quality factor $\tau\omega$, from fits to both the observed (solid symbols) quench dynamics in Fig. 6.7 and simulated (open symbols) dynamics using the LDA approach, as discussed in Section 6.1. The colour and symbol shape represents the quench direction, with $1/(k_F a_i) < 1/(k_F a_f)$ shown in red (point) and $1/(k_F a_i) > 1/(k_F a_f)$ shown in blue (diamond). Both axes are normalised by the Fermi energy $E_F = \hbar\omega_F = \hbar^2 k_F^2/(2m)$ determined from the density at the trap centre (refer to Section 2.4). **Left:** Observed and simulated (LDA) oscillation frequencies. Errorbars indicate the standard error from the fit, and the grey band represents the uncertainty in the simulation due to parameter variation. The dashed line shows the expected Higgs frequency $2\Delta_0$ at the maximum gap value Δ_0 at the trap centre (LDA). **Right:** Damping time for both simulated and measured dynamics, with errorbars denoting the propagated standard error. Data adapted from [135].

First, we note that the observed Higgs frequency increases with an increasing interaction parameter $1/(k_F a_f)$, as expected. However, we find that the measured oscillation frequencies ω (solid symbols) deviate by more than a factor of two from the frequencies predicted by

the LDA model (open symbols). Yet, deviations to the Higgs mode frequency derived from the density at the trap centre (dashed line) would be even greater, suggesting that the inhomogeneity of the samples indeed reduces the observed Higgs mode oscillation frequency.

Nonetheless, the measurement closely aligns with the LDA data in terms of the damping quality factor $\omega\tau$. We attribute this strong agreement in the quality factor to dephasing caused by the inhomogeneity of the samples [75, 215].

The following part of this section provides a more detailed discussion regarding the observed frequency deviations.

Deviations in frequency

Despite the good agreement in the damping quality factor $\tau\omega$, the central frequency shows significant deviations from the LDA model. The overlap between measurements and the LDA model depends on several parameters within the LDA model and the performed quenches that merit discussion.

First, the LDA model does not account for the fact that the order parameter Δ_∞ after the quench may differ from the equilibrium order parameter at the final interaction parameter $\Delta_\infty \neq \Delta_f$. As shown in [77, 79], the order parameter after the quench deviates more strongly from the equilibrium value as the quench size increases. To examine this effect, we have performed quenches with varying initial order parameters Δ_i and a fixed final order parameter Δ_f at unitarity, as presented in Fig. 6.5. However, no change in the measured oscillation frequency with varying quench size has been observed.

Second, deviations in the extracted frequency could be influenced by the determination of the Fermi energy ω_F , used here as a normalisation factor. As previously discussed, the Fermi energy is derived from the density $n(r=0)$ at the trap centre within LDA (refer to Section 2.4). This density is determined using the inverse Abel transformation, as explained in Section 4.2.1. As discussed there, deviations from an elliptical trapping potential may affect the reconstructed density profile, altering the determined Fermi energy. A more direct approach to check for a correct reconstruction of the density is to compare the atom number from the reconstructed in-situ density profiles to the atom number from more dilute absorption images after additional expansion time in TOF. We estimate the overestimation of the density to be below a factor of two [135]. Due to the scaling of the Fermi energy with density $E_F \propto n^{2/3}$ (see Section 2.4), the maximum discrepancy can not account for a shift of more than 60% in observed frequency. Thus, it can be ruled out that a systematic error in the reconstruction of the density is responsible for the large observed deviations in frequency.

Finally, the LDA model itself may be subject to significant variance with the given input parameters. As discussed in Section 6.1, the LDA model relies on assumptions about the Higgs profile, including temperature, linewidth and the dependence of the density of oscillators with the order parameter Δ . Variations in these parameters indeed affect the simulated oscillation frequency in the LDA model. The range of frequencies from these parameter variations is represented by the grey band in Fig. 6.8. However, the uncertainty due to parameter variation does not fully account for the observed discrepancy between measurements and the LDA model.

While each of these factors can explain some of the observed deviation, none sufficiently account for the large discrepancy seen in the data. Even if we combine the maximum

variation from all three factors, only extreme variations would bring the predicted frequencies closer to the observed data. Understanding the range of possible deviations from all three factors as probability distribution, this would yield an extremely low probability to hold true. Consequently, the deviations of the Higgs mode frequency in an inhomogeneous trapping potential can not be explained solely by the given LDA model.

6.2.4 Conclusion

In conclusion, our quenches probe region II of decaying oscillations in the quench phase diagram shown in Fig. 6.5. We observe oscillations across a wide range of quenches and can extract the oscillation frequency from phenomenological fits to the observed time dynamics. For quenches approaching the BEC regime, we observe fewer discernible oscillations, and for quenches furthest on the BEC side, oscillations become undetectable, signalling the vanishing of particle-hole symmetry as the BEC limit is approached.

For a quantitative comparison, we employ an LDA model for the Higgs spectrum in an inhomogeneous trapping potential. While there is striking agreement in the damping constant as reflected in the quality factor, the central frequency shows a deviation by more than a factor of two. We identify several parameters that could account for this discrepancy, but none fully bring the simulation fully into alignment with the measurements. The strong agreement in the quality factor suggests that the accelerated decay can be attributed to dephasing caused by the inhomogeneous distribution of local Higgs mode oscillators in an LDA framework. Hence, additional research is necessary to gain a more comprehensive understanding of the Higgs mode behaviour in an inhomogeneous trap.

6.3 Parametric excitations of a superfluid - Spectroscopic measurement of the Higgs mode

This section presents the results of the spectroscopic probing of the Higgs mode via a modulation of the magnetic field. As previously mentioned, this project has mainly been conducted by Andreas Kell [134], while my contributions include data acquisition, discussions and interpretation of the data. A more detailed discussion about the experiment, methodology and findings is available in Andreas Kell's PhD thesis [134]. Here, we provide a brief overview of the experimental aspects and the analysis of the obtained results. Very recently, the findings have been published in [100].

In earlier works [218], parametric excitations via magnetic field modulations have been used to probe the excitation spectrum of a Fermi gas in the BEC-BCS crossover. Before that, the excitation spectrum of a strongly interacting Fermi gas has been studied using radio-frequency-transitions (RF) [219]. A key advantage of magnetic field modulations over RF-transitions is that they do not involve a third state, making the process of probing more straightforward and simplifying the observations.

In previous experiments, RF-transitions have been used by my predecessors to excite the Higgs mode in our experiment [78]. RF-transitions couple to the system by addressing the density [40, 82], using Rabi oscillation to couple to a third unoccupied state. This approach, however, affects momentum states differently, leading to a broad excitation spectrum.

An alternative approach to drive parametric excitations is to couple directly to the interaction parameter $1/(k_F a)$ via the scattering length a [79, 218]. To this end, the magnetic field is modulated, resulting in an effective modulation of the scattering length and interaction parameter through the associated Feshbach resonance (refer to Section 2.2.2). This method has been employed previously to dissociate composite dimers on the BEC side of the BEC-BCS crossover by modulating the field at a frequency exceeding the corresponding two-body bound state energy E_B [218]. We use this method again in Chapter 7 to selectively break composite dimers.

6.3.1 Experimental realisation

In this experiment, we induce parametric oscillations of the interaction parameter by modulating the magnetic field. For this, we again use the custom-built small magnetic field coil (refer to Section 4.5.2) to superimpose a rapidly changing magnetic field on top of the large offset magnetic field generated by the Feshbach coils (refer to Section 3.2.6). In contrast to the quench protocol in Section 6.2.1, the driving circuit of the small magnetic field coil is replaced by an arbitrary function generator and an amplifier circuit, as detailed in Section 4.5.2.

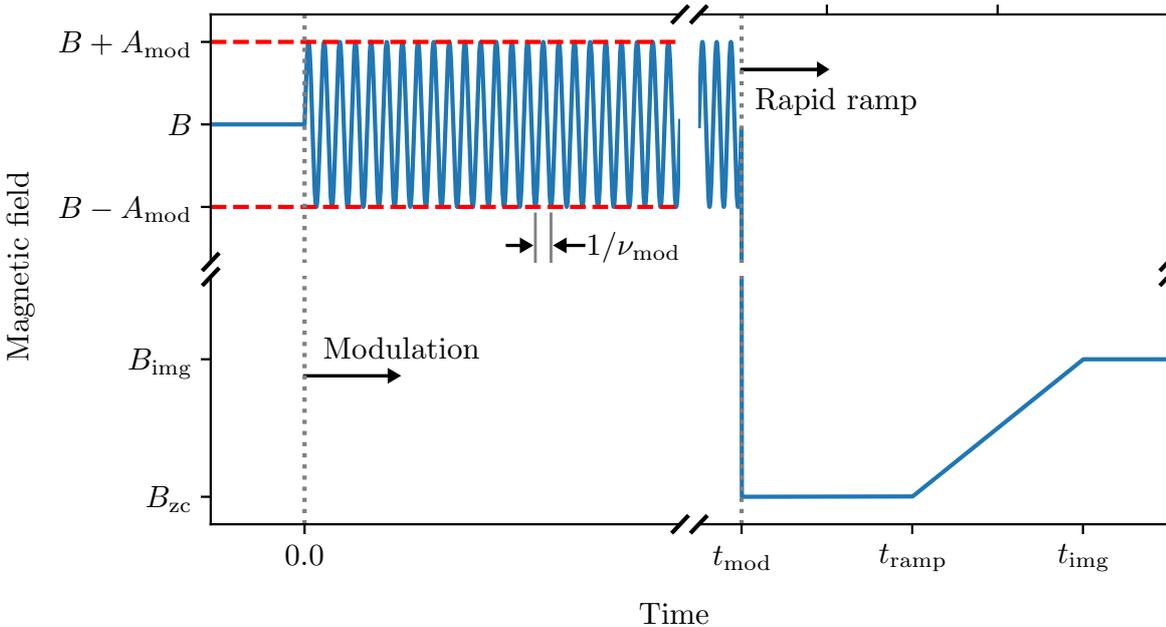


Figure 6.9: Experimental protocol to drive modulations. The figure shows the magnetic field profile used to probe the Higgs mode via a modulation of the interaction parameter $1/(k_F a)$. First, the sample is prepared at an offset field of B . At $t = 0$, the offset field is added with a sinusoidal modulated magnetic field with amplitude A_{mod} and frequency ν_{mod} for a duration t_{mod} using the small magnetic field coil, introduced in Section 4.5. The modulation time t_{mod} is rounded such that the modulation contains an integer number $n_{\text{mod}} = 2\pi t_{\text{mod}} \nu_{\text{mod}}$ of periods. After the modulation, the trap is immediately turned off at t_{hold} and the RR is initiated to measure the condensate fraction CF (see Section 4.3). The rapid ramp incorporates two timescales, namely the time $t_{\text{ramp}} - t_{\text{hold}}$ at which the sample is freely expanding at the zero crossing of the Feshbach resonance B_{zc} , as well as the ramp time $t_{\text{img}} - t_{\text{ramp}}$ to the final imaging magnetic field B_{img} . Figure adapted from [100].

The evolution of the total magnetic field profile used to drive parametric excitations in this experiment is shown in Fig. 6.9. Initially, the sample is prepared at a magnetic field B , created solely by the Feshbach coils. After preparation, at $t = 0$, the custom-built small magnetic field coil is activated to superimpose a modulation of the magnetic field. The settings of the small magnetic field coil comprises frequency ν_{mod} , amplitude A_{mod} and the modulation time t_{mod} . To maintain continuity in the magnetic field, the modulation time t_{mod} is rounded such that only an integer number of periods $n_{\text{mod}} = 2\pi t_{\text{mod}}\nu_{\text{mod}}$ is allowed. This is crucial to prevent unwanted quenches in the experiment.

After the modulation, the small magnetic field coil is switched off at t_{mod} , and immediately afterwards, the rapid ramp is initiated, as described in Section 4.3. This protocol mimics the proposed *protocol B* from [79]. According to the authors, the ramp time of the rapid ramp must not be faster than the timescale of the parametric excitation in order to observe Higgs mode excitations effectively.

6.3.2 Observed excitation rate

In this experiment, parametric excitations with a large variation in the amplitude of the scattering length a are favoured. Therefore, the sample is prepared in a $|13\rangle$ mixture, which features the narrower of the broad Feshbach resonances (refer to Section 2.2.2). This setup allows for the largest variation in the scattering length a under a modulation of the magnetic field B .

To probe the excitation spectrum, the decay rate of the condensate fraction CF is measured following the modulation. In general, we can vary the modulation amplitude A_{mod} , modulation frequency ν_{mod} as well as the modulation time t_{mod} . We have observed that for each frequency, the condensate fraction CF decays exponentially with increasing modulation time t_{mod} or an increase in the square of the modulation amplitude A_{mod} , expressed by

$$CF = CF_0 \left(-\Gamma_i A_{\text{mod}}^2 t_{\text{mod}} \right), \quad (6.9)$$

where Γ_i is the frequency-dependent decay rate. This behaviour aligns with the previous experimental observations in dimer dissociation spectra using magnetic field modulations [218] and theoretical considerations of the excitation spectrum in [220]. To avoid introducing a new timescale and causing additional thermal decay of the condensate, we keep the modulation time constant at $t_{\text{mod}} = 100$ ms throughout all measurements, adjusting only the modulation amplitude A_{mod} and modulation frequency ν_{mod} .

For each measurement, the magnetic field modulation follows the protocol sketched in Fig. 6.9. In Fig. 6.10, the decay of the condensate fraction CF is shown for various modulation frequencies ν_{mod} and increasing modulation amplitudes A_{mod} for an exemplary sample at unitarity. We observe that for each frequency, the condensate fraction CF decays exponentially as per Eq. (6.9), with a decay rate of Γ_i that varies with frequency.

For stronger modulation amplitudes A_{mod} , the decay deviates from the simple model in Eq. (6.9). To account for larger modulation amplitudes A_{mod} , the model is extended by a second-order term to

$$CF = CF_0 \exp \left(-\Gamma_i A_{\text{mod}}^2 t_{\text{mod}} - \gamma \left[A_{\text{mod}}^2 t_{\text{mod}} \right]^2 \right), \quad (6.10)$$

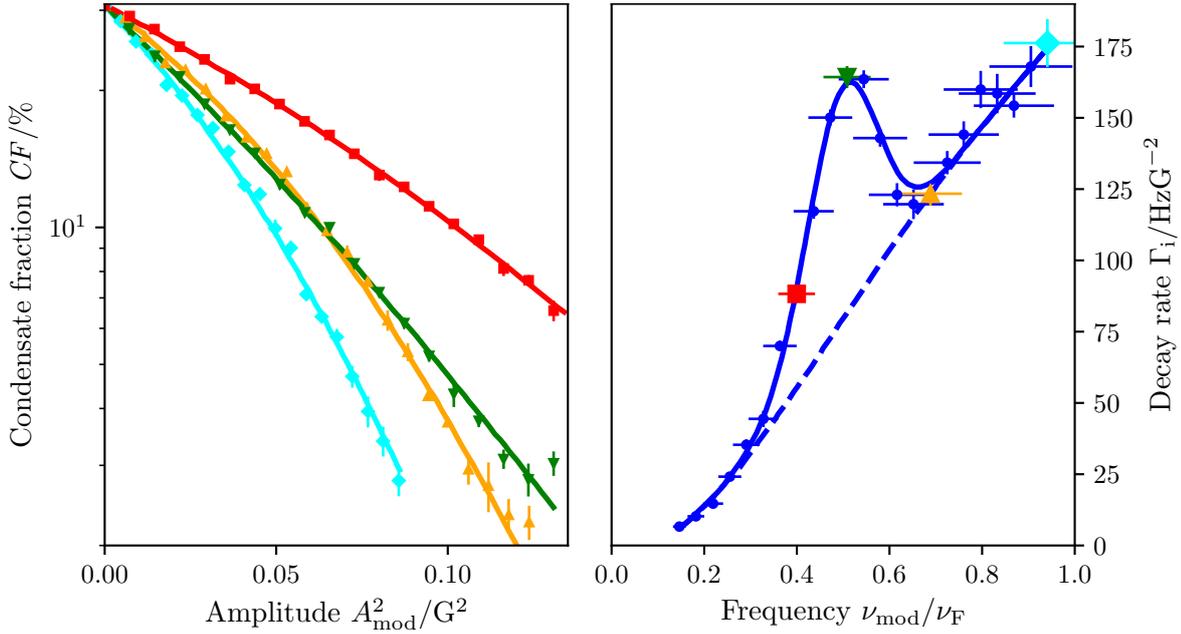


Figure 6.10: Decay of the condensate fraction after parametric excitation for a sample prepared at unitarity. **Left:** Measured condensate fraction CF after excitation via the modulation of the magnetic field with an amplitude of A_{mod} , as sketched in Fig. 6.9. Different lines represent different excitation frequencies ν_{mod} of 6.8 kHz (red squares), 8.6 kHz (green downfacing triangles), 11.7 kHz (orange upfacing triangles) and 16.0 kHz (turquoise diamonds). **Right:** Each line in the left panel is fitted to the extended model in Eq. (6.10), yielding an initial decay rate Γ_i . The decay rate is plotted as a function of the modulation frequency ν_{mod} (blue points), with corresponding datapoints from the **left** panel highlighted by the same colours and symbols. The resulting decay rate profile is fit by an empirical function (blue line), incorporating a polynomial background term (dashed line) as defined in Eq. (6.11). Figure adapted from [100].

where γ is a correction parameter for the second-order term. This extended model allows for more robust fitting and prevents an overestimation of the initial decay rate Γ_i . A fit of this model to the respective data is shown in the left panel of Fig. 6.10.

After fitting the initial decay rate Γ_i across various frequencies at a given magnetic field, the frequency dependence of the initial decay rate is identified and plotted in the right panel of Fig. 6.10. We observe that the decay rate varies with frequency, displaying a prominent peaked feature on top of an increasing background. The background is caused by collective excitations, such as phonon excitations of the (Nambu-)Goldstone mode at low frequencies (refer to Section 2.3.4) and incoherent pair-breaking excitations at higher frequencies near the pairing gap Δ [221]. For more detail about the decay rate at low frequencies, consult the PhD thesis of Andreas Kell [134]. We model the background with a 4th degree polynomial function and the pronounced feature on top of the background by a Gaussian function, yielding the model function

$$\Gamma_i(\nu_{\text{mod}}) = \sum_{k=1}^4 a_i \nu_{\text{mod}}^k + A \exp\left(-\frac{(\nu_{\text{mod}} - \nu_0)^2}{2\Delta\nu^2}\right). \quad (6.11)$$

Here, the polynomial background is defined by the empirical factors a_i , while the Gaussian

function is characterised by an amplitude A , a central frequency ν_0 and a width $\Delta\nu$. We identify the symmetric peak with the collective Higgs mode excitation. A fit to the excitation spectrum of the initial decay rate Γ_i is shown in the right panel of Fig. 6.10. This yields the resonance frequency ν_0 and spectral width $\Delta\nu$ for a sample prepared at a specific interaction parameter $1/(k_F a)$ and temperature T .

Next, this analysis protocol will be employed for several measurements across the strongly interacting regime of the BEC-BCS crossover.

6.3.3 Excitation rate across the crossover

In the previous Section 6.3.2, the methodology to determine the excitation spectrum has been introduced. The frequency dependent decay rate Γ_i of a condensate is measured following a suitable excitation of the superfluid via a modulation of the magnetic field. This measurement is performed for a selected initial temperature T and interaction parameter $1/(k_F a)$, set by the magnetic field during preparation.

In order to probe the excitation spectrum across the crossover, we repeat the same experimental protocol (see Section 6.3.1) and analysis (see Section 6.3.2) for samples prepared at different magnetic fields, i.e. interaction parameters $1/(k_F a)$, and temperatures T . As introduced in Section 5.3.1, the temperature of the sample is tuned by a time-varied decompression and a subsequent recompression of the trapping potential. Note, that the temperature range in this experiment is constrained by the need for an observable condensate and sufficient signal-to-noise ratio in the decay rate. To improve the robustness of the fitting method at small condensate fractions, we simplify the fitting model in Eq. (6.10) by setting $\gamma = 0$, disabling the second-order correction. The excitation spectra across various interaction parameters and temperatures are shown in Fig. 6.11.

For all probed interaction parameters and temperatures, the excitation spectrum consistently exhibits a peaked feature on top of an increasing background, as modelled by Eq. (6.11). We use this model function to extract the centre position and width of the peaked Gaussian on top of the background. An overview of the extracted centre frequencies ν_0 and widths (FWHM) $\Delta\nu_{\text{FWHM}}$ as functions of the interaction parameter $1/(k_F a)$ and temperature T is shown in Fig. 6.12. A discussion about the determined results is given in the following parts of this section.

Observed resonance frequency

We first examine the measured centre frequencies ν_0 of the observed resonances. With an increasing interaction parameter $1/(k_F a)$, we observe an increase in the resonance frequency ν_0 . Both the frequency and the interaction parameter are normalised by the Fermi energy derived from the density at the trap centre, refer to Section 2.4. As discussed in Section 6.1, we have developed an LDA model to compare our observed spectra with theoretical predictions of the Higgs mode spectra in an inhomogeneous sample. Unlike the results in Section 6.2, the LDA model shows excellent agreement with the observed resonance positions.

We have also measured the excitation spectrum for various temperatures T/T_F for a gas prepared at unitarity $1/(k_F a) = 0.0$. Due to the need for a measurable condensate fraction and sufficient signal-to-noise ratio, the temperature range is limited to $T/T_F \sim 0.15$, as

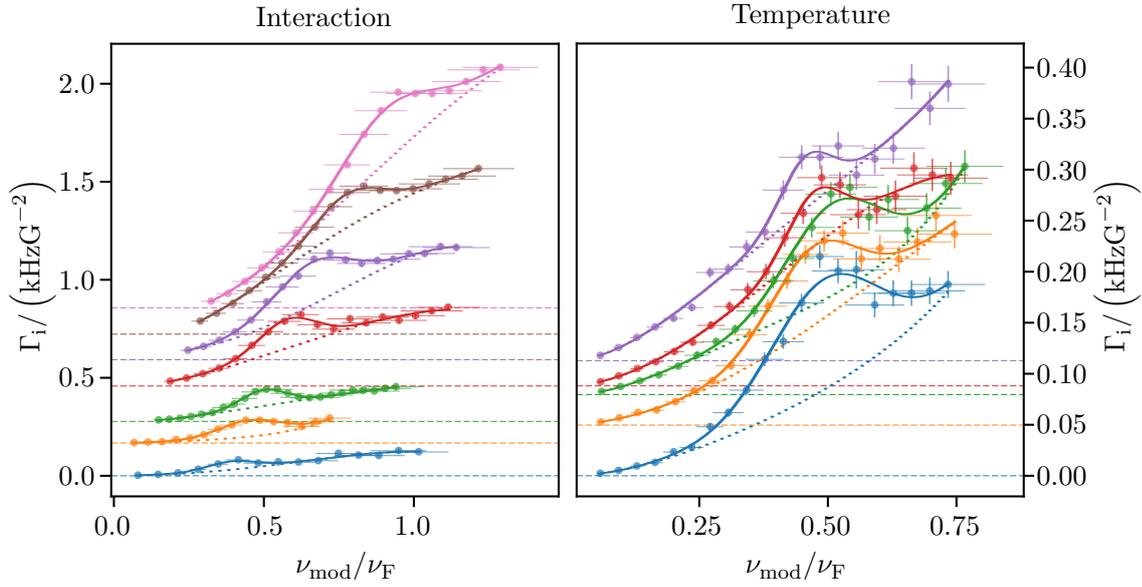


Figure 6.11: Measured excitation spectra for various initial interaction parameters and temperatures. **Left:** Initial excitation rate Γ_i (points) as a function of the modulation frequency ν_{mod} with fits (lines) to the model function in Eq. (6.11), incorporating a polynomial background function (dashed lines). Each colour represents a different initial interaction parameter $1/(k_{\text{F}}a)$, which are (top to bottom) 0.39, 0.30, 0.21, 0.12, 0.0, -0.07, -0.18. For clarity, data are vertically shifted with an offset of $(1.5/(k_{\text{F}}a) + 0.28) \text{ kHzG}^{-2}$. **Right:** Initial decay rate Γ_i for a sample prepared at unitarity but for different initial temperatures T/T_{F} . Here, the colour denotes the temperature T/T_{F} , which is (bottom to top) ranging from 0.08 to 0.13. Again, data are vertically shifted by an artificial offset of $2(T/T_{\text{F}} - 0.08) \text{ kHzG}^{-2}$. Horizontal errorbars indicate the standard error on ν_{F} , while the vertical errorbars show the statistical standard error from the fit. Figure adapted from [100].

indicated by the grey-shaded area. Contrary to the expected temperature dependency in Fig. 6.2, we observe no significant correlation between resonance frequency and temperature. This result, however, aligns with the recent measurement of the Higgs mode using Bragg spectroscopy in [75]. Unlike our method, their approach probes a quasi-homogeneous region in the centre of the inhomogeneous sample by focused two photon Bragg spectroscopy following a quench of the interaction potential. While observing no temperature dependency as well, their observed centre frequencies generally lie higher, as they correspond to the (homogeneous) Higgs mode frequency at the centre of the trap.

Generally, the missing temperature dependency could indicate the existence of a pseudogap state above the superfluid critical temperature [222, 223]. However, further research is needed to understand the exact mechanism behind this observation.

Observed resonance width

The second parameter extracted from the fits to Eq. (6.11) is the resonance width. For all plots, we convert the width of the Gaussian model $\Delta\nu$ to the full width at half maximum

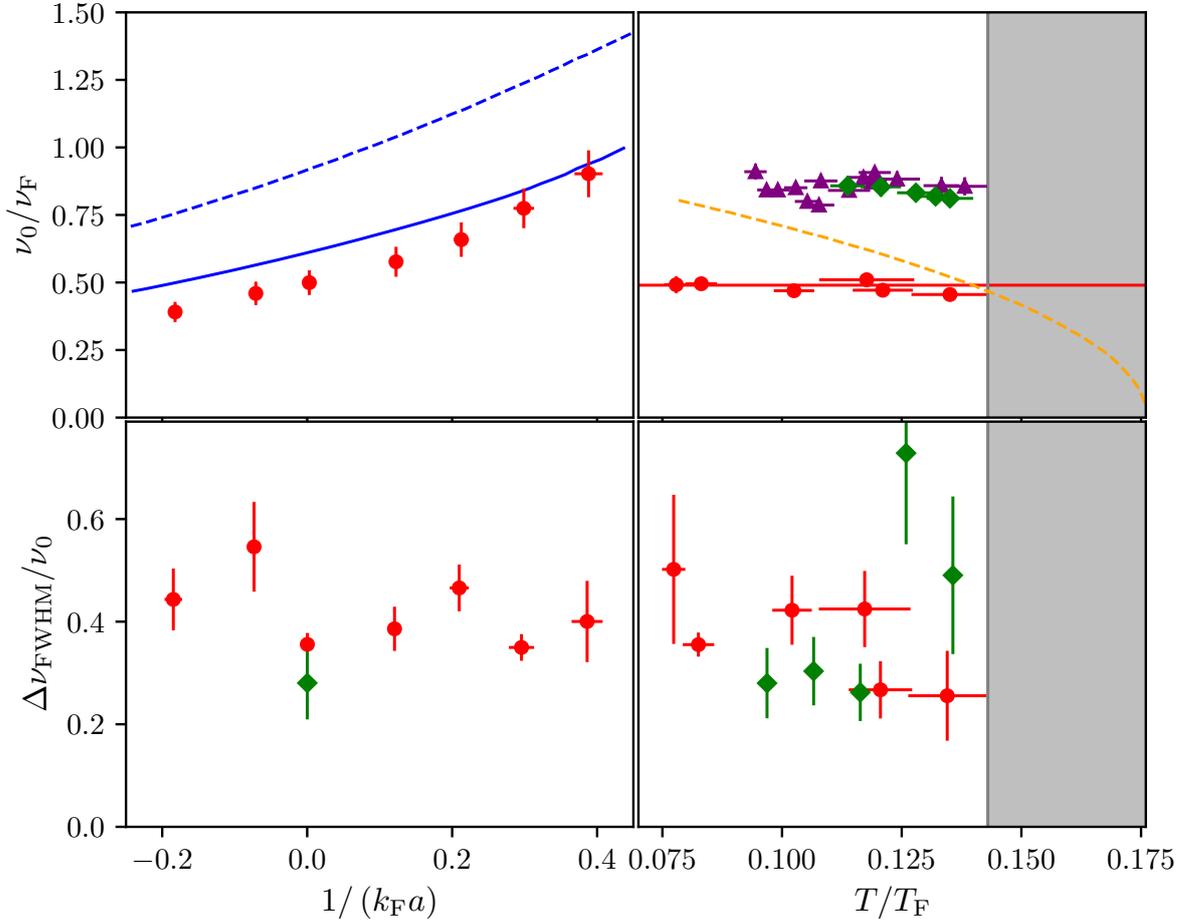
(FWHM) $\Delta\nu_{\text{FWHM}}$.


Figure 6.12: Extracted resonance positions and widths for various initial interaction parameters and temperatures. The figure presents the extracted resonance positions ν_0 and widths (FWHM) $\Delta\nu_{\text{FWHM}}$ as functions of the initial interaction parameter $1/(k_F a)$ and temperature T . **Upper left:** Measured (red points) resonance positions ν_0 are plotted against $1/(k_F a)$. The solid (blue) line shows the LDA simulation results from Section 6.1, while the dashed (blue) line corresponds to the Higgs mode frequency derived from the density at the trap centre. **Upper right:** Measured resonance positions (at unitarity) are plotted as a function of the temperature T/T_F . No significant temperature dependence is observed, and the data is well represented by their average (red solid line), contrasting with the expected temperature dependence in Fig. 6.2. Our data is compared to a different recent study in [75] (green diamonds and purple triangles). **Lower left:** Resonance width (FWHM) $\Delta\nu_{\text{FWHM}}$ as a function of $1/(k_F a)$, showing minimal dependence. Again, green diamonds are from [75]. **Lower right:** Resonance width (FWHM) $\Delta\nu_{\text{FWHM}}$ as a function of T/T_F , showing a weak dependence at most. Again, green diamonds are from [75]. The grey area marks the region of insufficient signal-to-noise ratio. Errorbars include the systematic uncertainty in the Fermi energy E_F (refer to Section 4.2.1) and statistical standard errors from the fit. Figure adapted from [100].

Overall, we observe an almost constant width across the crossover for a varying interaction parameter $1/(k_F a)$, or for a varying temperature T at unitarity. Our data suggests a weak dependency of the resonance widths with temperature at most. Interestingly, the recent

observation of the Higgs mode in [75] reports a spectral width similar to ours, which is notable given that they probe a quasi-homogeneous sample, while our experiment is examined in an inhomogeneous sample.

In Section 6.1, we have developed an LDA model to incorporate dephasing of an inhomogeneous collection of Higgs oscillators, making up the total spectrum. An exemplary spectrum from this model is shown in Fig. 6.3. However, this model predicts spectra that are significantly broader (by a factor of $\sim \times 2$) than those observed in our experiment in Fig. 6.12.

Both the agreement with the quasi-homogeneous measurement of the Higgs mode from [75] and the strong deviations from the LDA model suggest that the spectral width in an inhomogeneous trap is not dominated by dephasing as predicted by the LDA model.

6.3.4 Conclusion

In conclusion, we have employed a parametric excitation scheme via the modulation of the magnetic field to probe the excitation spectrum of a superfluid Fermi gas and to detect the Higgs mode. The observed Higgs resonance frequencies reveal excellent agreement with the simulated spectra in an inhomogeneous trapping potential using the LDA approach (see Section 6.1). However, the LDA approach significantly overestimates the observed resonance width, suggesting that the spectral width is not dominated by the effect of dephasing.

Moreover, our measurements indicate no observable temperature dependence of the resonance frequency. This observation aligns with a recent measurement of the Higgs mode [75], but contrasts with existing theoretical predictions in Fig. 6.2, highlighting a need for further research.

6.4 Comparison between the excitation techniques

In this chapter, we have explored the Higgs mode using two distinct excitation techniques: quenches and parametric excitations of the interaction parameter $1/(k_F a)$. Broadly speaking, quenches allow for the Higgs mode to be induced and observed in a time-resolved manner by measuring oscillations in the condensate fraction CF – effectively a time-domain measurement of the Higgs mode. In contrast, parametric excitations serve as a spectroscopic approach of measuring the Higgs mode in the frequency domain, where the mode is observed through an enhanced decay of CF for a modulation with a frequency near the Higgs mode frequency.

In the experiments in Section 6.2 and Section 6.3, the excitation techniques leverage the coupling between the magnetic field B and the scattering length a (see Section 2.2.2). This coupling enables us to map variation in the magnetic field to changes in the scattering length a , allowing us to excite the Higgs mode via either a quench or a modulation of a superimposed magnetic field, generated by the custom-build small magnetic field coil. Using this approach, we have investigated the Higgs mode in a strongly interacting Fermi gas across the BEC-BCS crossover, confined in a $|13\rangle$ mixture within a harmonic, inhomogeneous dipole trap potential.

Due to the trap's inhomogeneity, the expected Higgs becomes more complex. Here, we have developed an LDA model of the expected Higgs mode spectrum, that accounts for the sample's inhomogeneity and, by performing a Fourier transformation on the spectrum, derived anticipated Higgs mode oscillations. Together, these models provide predictions for the role of inhomogeneity on both the time- and frequency-domain measurements.

In both measurements, we have observed a Higgs mode frequency significantly lower than the expected (homogeneous) Higgs mode frequency derived from the density at the centre of the trap, as anticipated. For parametric excitations, the reduced central frequency predicted by the LDA model aligns well with the observed central frequency. However, for the quench technique, we have found a notable discrepancy: the observed frequency was much lower than the LDA prediction, with the model's frequency exceeding the experimental value by more than a factor of two. Despite exploring possible explanations, we were unable to reconcile the model with the observed results.

Damping of the Higgs mode presents a second key observable, expected to be broadened by dephasing due to trap's inhomogeneity. For the quench measurement, the quality factors of the observed oscillations closely match those predicted by the LDA model, suggesting that the observed decay results from dephasing among oscillators at varying Higgs frequencies within the inhomogeneous sample. In contrast, for parametric excitations, we found a nearly constant resonance width across the strongly interacting regime and for different temperatures at unitarity. Remarkably, the observed resonance width in our inhomogeneous trap closely matches that of an effectively homogeneous system from a different recent study [75]. Furthermore, the LDA model overestimates the expected width by a factor of ~ 2 compared to our parametric excitation measurement. This indicates that the parametric excitation technique observes a resonance width more akin to a homogeneous sample, less affected by dephasing due to inhomogeneities.

The differences between the results from quenching and parametric excitations can not be explained by the given model, indicating the need for further research to fully understand how these discrepancies are caused by the inhomogeneity of the trapping potential.

Cooling of a strongly interacting quantum gas by dimer dissociation

This chapter presents the mechanism and measured performance of the novel cooling technique developed in this thesis. Briefly, this new cooling technique leverages a shift in the mean-field energy during dissociation of composite dimers, enabling the selective removal of high-energy atoms from an inhomogeneous trap and thereby cooling the remaining sample.

As discussed in Section 2.3.1, a strongly interacting Fermi gas of two distinct spin states supports a two-body bound state for positive scattering lengths $a > 0$, enabling the formation of composite bosonic dimers with bound state energy E_B . Due to the Pauli exclusion principle, these composite bosons exhibit a reduced scattering length with other bosons and unpaired fermionic atoms, resulting in a weakened mean-field (repulsive) energy. To dissociate dimers, we apply a magnetic field modulation at a frequency ν_{mod} greater than the bound state energy E_B/h .

In first order, each dimer in the inhomogeneous trap has the same dissociation threshold frequency of $\nu_{\text{mod}} > E_B/h$. However, in second order, the dissociated fermionic fragments experience a greater mean-field energy within the remaining sample, causing an increase in the dissociation threshold. The shift $\Delta\nu_{\text{cool}}$ in dissociation threshold from the pure bound state energy E_B is due to the shift in mean-field energy from the dimer-dimer state to twice the atom-dimer state, with the shift being proportional to the inhomogeneous density $\Delta\nu_{\text{cool}} \propto n$ [108, 224, 225].

By breaking a dimer, the energy surplus of the modulation $E_B + \Delta\nu_{\text{cool}} - \nu_{\text{mod}}$, is converted into the kinetic energy of the fragments, potentially expelling them from the trap. The shift $\Delta\nu_{\text{cool}}$ thus underpins this new cooling technique, enabling selective dissociation and removal of high-energy atoms from the edges of the inhomogeneous trap, where the density approaches zero.

This approach allows for a cooling protocol akin to standard radio-frequency-transition (RF) evaporative cooling in a magnetic trap, adapted here for a spin mixture in an inhomogeneous optical dipole trap. The objective of any evaporative cooling scheme is to selectively remove the highest energy atoms from a sample, reducing its temperature after brief re-thermalisation. This concept is illustrated in Fig. 7.1.

Initially, the sample contains N_i atoms at temperature T_i with phase-space density $\tilde{\rho}_i$. By

removing $-\Delta N$ of the highest energy atoms, the average energy in the sample decreases, resulting in a temperature reduction to $T_f < T_i$. When performed efficiently, the phase-space density $\tilde{\rho}_f$ increases while the atom number decreases, leading to a drop in entropy.

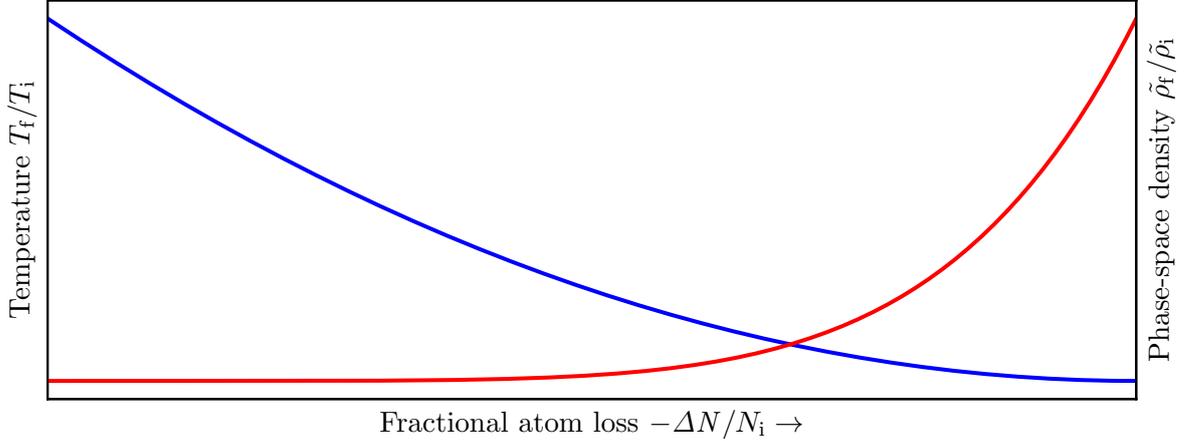


Figure 7.1: Goal of evaporative cooling. The figure illustrates the desired effect of evaporative cooling. Initially, a sample of N_i atoms is prepared at a temperature T_i with a phase-space density of $\tilde{\rho}_i$. By removing $-\Delta N$ atoms with higher-than-average energies, the temperature reduces to $T_f < T_i$ (blue line) and the phase-space density increases to $\tilde{\rho}_f > \tilde{\rho}_i$ (red line), achieving cooling.

Compared to standard cooling methods like RF-transition evaporation in a magnetic trap or forced evaporation in an optical dipole trap near a Feshbach resonance, this technique offers two main advantages. First, dissociation via magnetic field modulation avoids coupling to a third state, unlike standard RF-transitions, which helps to prevent strong three-body losses. Second, selective dissociation via the modulation of the magnetic field requires no trap alterations, preserving the depth of the trapping potential and making it easier to remain in the efficient runaway regime – meaning an increasing collision rate due to increasing density [85] – sustaining rapid thermalisation. This novel technique achieves high cooling efficiencies that meet and even surpass the highest reported efficiencies near a Feshbach resonance [85, 87, 91–99].

This work presents the second main result of this thesis. A manuscript covering this work is currently under review, with a preprint available in [102].

Outline of this chapter

The sections of this chapter are outlined as follows:

- Section 7.1: In the first section, we introduce the concept of evaporative cooling and discuss established state-of-the-art implementations.
- Section 7.2: Our newly developed evaporative cooling technique works by a modulation of a superimposed magnetic field and complements established methods. This section presents our technique and explains its theoretical basis.

- Section 7.3: In this section, we discuss the experimental implementation and show first results of the cooling technique in our system, focusing on the effect of the modulation parameters – frequency ν_{mod} and amplitude A_{mod} – on the cooling efficiency.
- Section 7.4: The cooling efficiency is a key parameter for any evaporative cooling technique. Here, we quantify our cooling technique and measure its efficiency for various magnetic fields, as well as in a gravity-tilted and gravity-compensated trap.
- Section 7.5: Another important aspect of any cooling technique is the intrinsic timescale required to cool the sample. In this section, we estimate the shortest feasible time for effectively applying our cooling technique.
- Section 7.6: Finally, we summarise and conclude the findings.

7.1 Principle of evaporative cooling

Evaporative cooling encompasses all techniques that reduce a sample's entropy by removing or isolating its high-energy components. In thermal equilibrium, a sample's state is governed by its statistical distribution: the Boltzmann distribution in Eq. (2.1) for high temperatures, or Bose and Fermi statistics in Eq. (2.5) and Eq. (2.2) for bosons and fermions, respectively. While these statistics differ conceptually, all show a higher likelihood of finding low-energy components, with the probability of high-energy components decaying in a manner similar to the exponential Boltzmann distribution. Consequently, the Boltzmann distribution often suffices to describe the removal of high-energy components from the distribution's tail.

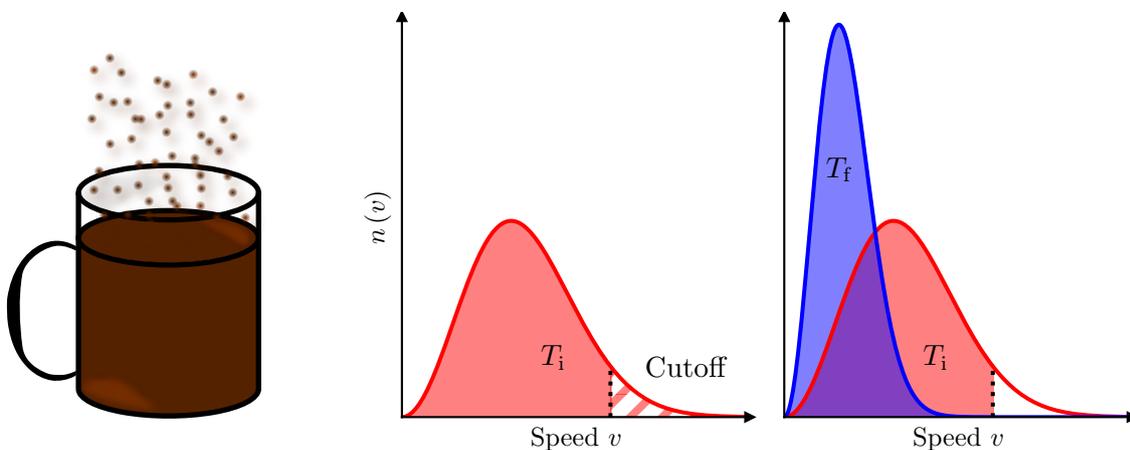


Figure 7.2: Principle of evaporative cooling. In a thermal sample, such as a cup of coffee or an atomic gas, the particles' energy distribution is governed by the (Maxwell-)Boltzmann distribution. Initially, at high temperature T_i , the sample contains many high-energy or high-speed v particles. By selectively removing particles with the highest energy above a cutoff, more energy per particle is removed from the sample than is available on average. Following this removal and a sufficient re-thermalisation period, the particles redistribute into a narrower distribution with a lower temperature $T_f < T_i$, achieving cooling. Analogously, for a cup of coffee, particles with high enough energy to enter the gaseous phase leave the liquid, resulting in effective cooling.

In Fig. 7.2, we illustrate the principal of evaporative cooling using the familiar example of a hot cup of coffee. Initially, the state of the coffee, i.e. the energy distribution among its particles, can be modelled by a Boltzmann distribution with an initial temperature T_i . Here, we introduce the Maxwell-Boltzmann distribution $n(v)$, which integrates the Boltzmann distribution across the full phase-space except for the particles' speed v [170]. Only particles exceeding a threshold speed v_{gas} (the cutoff) can transition from the liquid to the gaseous phase. In this phase, particles may escape the sample, effectively cooling it.

The temperature reduction occurs in two stages. First, high energy particles are expelled, thus removing them from the high-energy tail of the (Maxwell-)Boltzmann distribution. This truncated distribution no longer reflects the system's equilibrium state. After a suitable re-thermalisation period, the remaining particles rearrange into a narrower (Maxwell-)Boltzmann distribution with a lower temperature $T_f < T_i$, establishing the new equilibrium state.

This is the fundamental concept of evaporative cooling, applicable to many systems. The primary requirement for efficient evaporative cooling is the selective removal of high-energy atoms from the distribution.

7.1.1 Established evaporative cooling techniques

In ultracold atomic gases, evaporative cooling has been instrumental in reaching extremely low temperatures, enabling the experimental observation of highly sought-after quantum states. These states include the first observation of Bose-Einstein condensation (BEC) [30], degenerate Fermi gases [31] and the unitary Fermi gas [38, 39]. For all of these experiments, atomic samples were pre-cooled through successive stages to achieve temperatures necessary for observing the desired states. Pre-cooling stages utilise standard laser cooling techniques [25], after which the samples are trapped and further cooled via evaporative cooling in magnetic or optical dipole traps. Both of these methods are illustrated in Fig. 7.3, and are explained in more detail below.

RF-transition cooling in magnetic traps

The first observation of BEC in bosonic ^{87}Rb was made possible through RF-transition evaporative cooling in a magnetic trap, as illustrated on the left side of Fig. 7.3.

As discussed for our experiment setup in Section 3.2.5, atoms are trapped in a magnetic field by being pumped into a low-field-seeking state $|\uparrow\rangle$. Atoms in a state with a flipped magnetic moment (high-field-seeking state $|\downarrow\rangle$) are anti-trapped, causing them to be expelled from the magnetic trap. Due to the Zeeman shift explained in Fig. 3.2, an atom's energy depends on the magnetic field, and therefore, on its position within the magnetic trap. This shift corresponds to a change in the RF-transition frequency required to transfer an atom from the trapped state $|\uparrow\rangle$ to the anti-trapped state $|\downarrow\rangle$. Atoms at the edges of the trap are subject to a stronger magnetic field and require a higher RF-transition frequency to be transferred to the anti-trapped state.

Because the trap edges are typically occupied by high-energy atoms, this frequency shift allows selective transfer of high-energy atoms to the anti-trapped state, thereby removing them from the sample.

In practice, RF-transition evaporative cooling is optimised by ramping from high to low

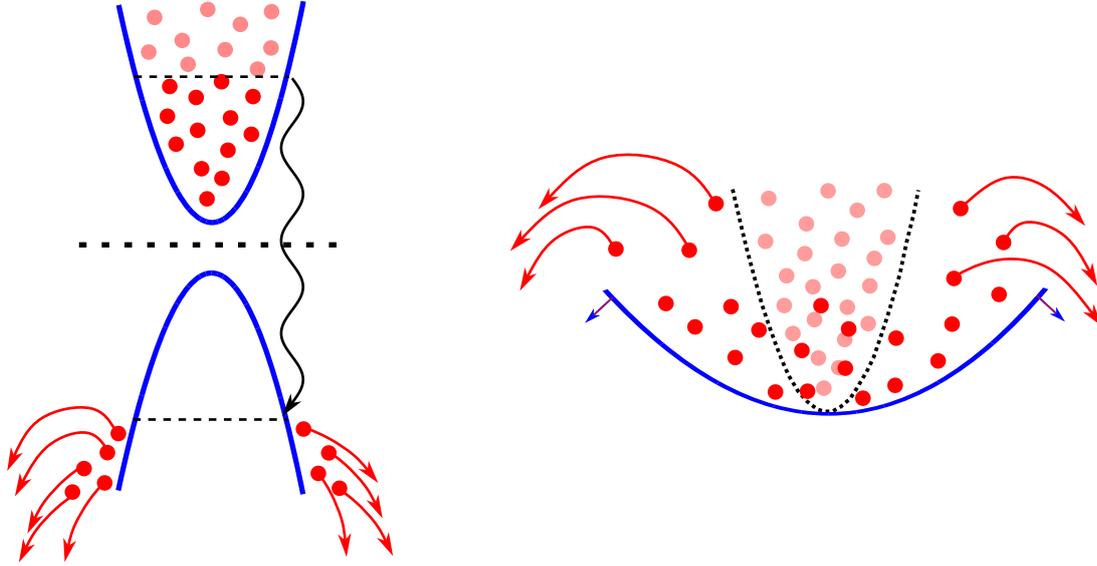


Figure 7.3: Established evaporative cooling techniques. Our new cooling technique complements the established methods of RF-transition cooling in a magnetic trap (**left**) and forced evaporation in an optical dipole trap near a Feshbach resonance (**right**). **Left:** In RF-transition cooling, the sample is prepared in a trapped state that is flipped to an anti-trapped state through rf-transitions. Due to the Zeeman effect (refer to Fig. 3.2), atoms with higher energy, primarily located at the trap edges, have a higher transition frequency, allowing for selective removal of high-energy atoms. **Right:** In a dipole trap, there is no such transition. Instead, high-energy atoms are removed by gradually lowering the trap depth. However, as this reduces the trap depth, it also decreases the density, making it challenging to remain in the efficient runaway regime.

transition frequencies while allowing for efficient re-thermalisation of the remaining sample. Notably, this cooling method does not alter the trap itself, allowing the atoms to remain in the efficient runaway regime without slowing down thermalisation [85]. An overview of the ramp implemented in the magnetic trap stage of our experiment is given in Section 3.2.5.

While highly effective for single-state bosonic species, this method has limitations tied to the physics of magnetic traps, which are not designed to hold arbitrary state mixtures. Consequently, and due to the Pauli exclusion principle, this technique does not achieve the temperature required for reaching a degenerate Fermi gas.

Forced evaporation in an optical dipole trap near a Feshbach resonance

To reach a degenerate Fermi gas, the concept of evaporative cooling has been adapted for optical dipole traps. Unlike magnetic traps, optical dipole traps can confine arbitrary state mixtures, making them ideal for hosting strongly interacting Fermi gases necessary to prepare the degenerate Fermi gas [31] and to achieve the unitary Fermi gas [38, 39].

In optical dipole traps, however, there is no straightforward mechanism to selectively expel high-energy atoms independently from alterations to the trap itself. Thus, we rely on lowering the trap’s potential walls to predominantly remove high-energy atoms. To enable sufficiently fast thermalisation, the dipole trap is gradually ramped down while the magnetic

field is tuned close to a Feshbach resonance, enhancing thermalisation rates. This approach is applied in the final preparation stage of our experiment, as discussed in Section 3.2.6.

Although conceptually simpler than RF-transition cooling in a magnetic trap, forced evaporation in an optical dipole trap has a significant drawback related to the control and reduced confinement of the trap. As the power of the optical dipole trap is lowered, confinement diminishes throughout the entire trapping region, effectively reducing the gas density and, consequently, extending the thermalisation time [85]. This presents a major obstacle in reaching high efficiencies.

Alternative cooling techniques

Both cooling methods introduced so far remain standard, state-of-the-art techniques in most ultracold atomic gas experiments. However, the search for improved cooling techniques is an active area of research aimed at accessing novel quantum states at even lower temperatures. These sought-after quantum states include unconventional superconductivity [88], topological quantum states [89], and spin liquids [90].

To achieve the required temperatures, alternative cooling methods have been proposed, such as separating high- and low-entropy regions by potential sculpting [226–228]. Related approaches have been experimentally implemented, including potential shaping in the Hubbard model [229], or by introducing alternating high- and low-entropy sites via a superlattice potential along one spatial direction [230].

7.2 Evaporative cooling via dissociation of composite dimers

In this section we introduce and discuss the concept behind the novel cooling technique developed in this thesis. The working principle of this technique is illustrated in Fig. 7.4.

Initially, the gas is prepared in a harmonic dipole trap, as described in Section 3.2.6 and sketched in **a** in Fig. 7.4. The gas is prepared on the BEC side of the crossover, where a two-body bound state with energy E_B exists, and the trap is predominantly occupied by composite bosonic dimers from the original fermionic state mixture.

To break a dimer, we apply a magnetic field modulation, as detailed in Section 7.3. The required energy to break a pair includes the bound state energy E_B plus a shift δE , which originates from a shift of the repulsive mean-field energy as a dimer in the original dimer gas breaks into two dissociated fermionic fragments. As covered in Section 2.3.1, the shift arises due to the change in the scattering length from $a_{dd} = 0.6a$ [121] for a dimer-dimer interaction to $a_{ad} = 1.18a$ for the interaction between a fermionic fragment and a dimer [120].

More quantitatively, the shift can be calculated by comparing the mean-field energy of a composite dimer in gas of dimers (refer to Eq. (2.23)) with twice the mean-field energy of a fermion in the same gas of dimers. The shift of the mean-field energy [108, 224] can be calculated and reads

$$\delta E = (3a_{ad} - a_{dd}) 2\pi\hbar^2 n_\sigma / m_{\text{Li}}, \quad (7.1)$$

where n_σ is the density per spin state, equivalent to the density of composite dimers in a fully paired sample. The expression in parentheses captures the physics of breaking a composite dimer into two unpaired fermions with stronger repulsive mean-field energy in the remaining

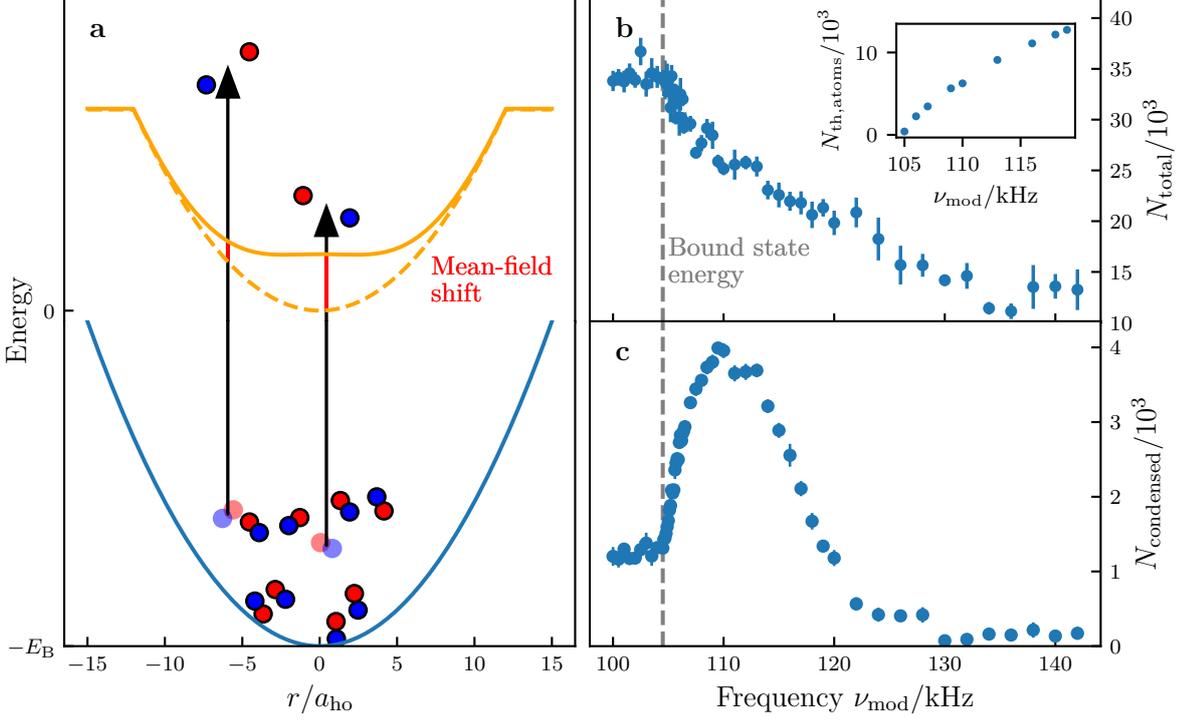


Figure 7.4: Working principle of cooling via dissociation of dimers. **a** Initially, atoms of the two spin states (blue and red circles) are confined within a harmonic dipole trap potential (blue line) and constitute a partially condensed quantum gas of composite dimers. The dissociation threshold frequency of the dimers depends on the two-body bound state energy E_B as well as the shift in mean-field energy from the initial dimer gas to the mixed gas of dimers and fermionic atoms, which originate as fragments of dissociated dimers. The shift of the mean-field energy (red line) is proportional to the density, resulting in a position-dependent dissociation threshold frequency that increases towards the trap centre. **b** This creates a sharp dissociation threshold frequency at the bound state energy E_B (dashed line) for atoms at the dilute edge of the trap. For higher frequencies, more dimers are dissociated and evaporated from the trap, as seen in the decreasing total atom number N_{total} . The increasing fraction of dissociated atoms is reflected by the growing number of thermal atoms $N_{th,atoms}$ (inset) during dissociation. **c** More notably, the condensed number of atoms $N_{condensed}$ exhibits a peaked structure. At frequencies close to E_B , evaporation selectively targets high-energy dimers at the dilute trap edges, increasing the number of condensed atoms. For higher frequencies, the evaporation becomes less selective, reducing the number of condensed atoms.

cloud of dimers. This expression involves a change of the reduced mass μ_r from $\mu_{r,dd} = m_{Li}$ (dimer-dimer) to $\mu_{r,ad} = 2m_{Li}/3$ (atom-dimer) and includes a factor of 2 to account for both fermions produced in dissociation, resulting in the shift of the mean-field energy given in Eq. (7.1).

The energy threshold for dimer dissociation is given by the sum of the bound state energy and the shift of the mean-field, which reads

$$E_{diss} = E_B + \delta E. \quad (7.2)$$

By modulating the magnetic field at a frequency ν_{mod} above E_{diss}/h , the fermionic fragments

gain kinetic energy, propelling them to higher trap states and potentially expelling them from the trap. This is illustrated in **a** of Fig. 7.4 and quantitatively shown for an exemplary gas in **b**, measured using the RR technique (see Section 4.3). For frequencies ν_{mod} above the bound state energy E_{B} , the total amount of remaining atoms (per spin state) N_{total} decreases with increasing frequency.* This has been measured after a modulation plus a brief thermalisation period. A similar trend is observed in the amount of dissociated atoms before thermalisation, indicated by number of thermal unpaired atoms $N_{\text{th,atoms}}$ (per spin state), which are probed by applying a π -pulse to transfer and image them in a third unoccupied spin state.

While the increase in dissociated atoms with frequency is evident, it is of particular significance that the mean-field shift enables the selective dissociation and removal of atoms at specific density shells. Dimers at the edge of the trap, where the density is lowest, have the highest energy. For these high-energy dimers, the threshold dissociation frequency approaches the pure bound state energy $E_{\text{diss}} \rightarrow E_{\text{B}}$. In contrast, low-energy dimers occupy the trap centre at much higher densities, requiring a higher dissociation threshold energy $E_{\text{B}} + \delta E$ to break them. The shifted dissociation threshold explains the increasing amount of dissociated atoms with frequency in **b** of Fig. 7.4, despite the reduced overlap between initial and final trap states for increasing excess energy [17]. More importantly, the shifted threshold protects the coldest dimers in the sample from direct dissociation when the frequency is near the bound state energy E_{B} , selectively dissociating only high-energy dimers.

To illustrate, for measured densities on the order of $10 \mu\text{m}^{-3}$ and calculated bound state energies around 100 kHz, the mean-field shift typically ranges from $\delta E = 1 \text{ kHz}$ to 10 kHz , depending heavily on the scattering length and position on the BEC side of the crossover.

By modulating the magnetic field at a frequency ν_{mod} close to the bound state energy E_{B} , primarily high-energy dimers at the trap edge are dissociated, reducing the remaining sample's energy per particle. In Fig. 7.4 **c**, this is observed as an increase in the number of condensed atoms $N_{\text{condensed}}$, indicating a rise in phase-space density and reduction in entropy. At frequencies significantly above the bound state energy, however, the number of condensed atoms decreases, eventually depleting the condensate, as the dissociation loses selectivity and dimers are broken throughout the trap, including the lowest energy dimers of the condensate.

Advantages over established methods

Our newly developed cooling technique complements the established methods discussed in Section 7.1.1. Compared to evaporative cooling via RF-transitions in a magnetic trap or forced evaporation in an optical dipole trap, this new method offers several major advantages.

First, this method enables cooling of an arbitrary two-spin mixture, provided the bound state energy on the molecular side of the Feshbach resonance is accessible. Second, the magnetic field modulation avoids coupling to a third spin state, thereby preventing strong three-body losses. Finally, the targeted and selective dissociation process achieved through magnetic field modulation keeps the trap unaltered throughout the evaporation process. This consistency aids in maintaining the efficient runaway regime, promoting rapid thermalisation since density does not decrease, unlike in forced evaporation within an optical dipole trap.

*Here, the modulation amplitude A_{mod} of the magnetic field modulation is chosen to optimise the signal. For further details on frequency and amplitude consult Section 7.3.

Later in Section 7.4, we analyse the efficiency of this cooling method, demonstrating that it reaches high efficiencies of $\gamma = 4$ in a gravity-tilted trap and $\gamma = 6$ in a gravity-compensated trap, comparable to and even surpassing the highest reported efficiencies in forced evaporation steps near a Feshbach resonance [85, 87, 91–99].

7.2.1 Evaporative cooling step model

To deepen our understanding of this cooling mechanism, we simulate the proposed evaporative cooling step using the step model of evaporation [85]. We simulate an interacting Bose gas with our measured trap frequencies, temperature and atom number. This defines the initial internal energy U by integrating the product of energy and occupation (given by the Bose distribution function in Eq. (2.5)) over the entire phase space, spanned by momentum $d^3\mathbf{p}$ and position $d^3\mathbf{r}$. After selectively removing $-dN$ high-energy atoms from the edges of the sample, we recalculate the internal energy to estimate the corresponding temperature change.

The step model is based on the proportionality between atom number N , temperature T and internal energy U , expressed as [85]

$$U = \xi N k_B T, \quad (7.3)$$

where ξ depends on the trap geometry, such as $\xi = 3/2$ for a homogeneous trap or $\xi = 3$ for a harmonic trap [8].

As discussed in Section 7.1, evaporative cooling removes a number $-dN = N_f - N_i$ of atoms, changing the initial atom number N_i to a final atom number of N_f . After removing these atoms and allowing for sufficient thermalisation, the remaining atoms rearrange into a modified (Boltzmann) distribution with a final temperature T_f . Thus, the sample's temperature changes by $dT = T_i - T_f$. Naturally, the remaining atoms also exhibit a change in internal energy $dU = U_i - U_f$, reflecting the change in atom number dN and temperature dT .

The step model of evaporation describes this cooling step as

$$\begin{aligned} U_f &= U_i - dU = \xi (N_i - dN) k_B (T_i - dT) \\ \Leftrightarrow U_f &= U_i - dU = \xi k_B (N_i T_i - N_i dT - T_i dN + \mathcal{O}(dN dT)), \end{aligned} \quad (7.4)$$

where the final non-linear term $\mathcal{O}(dN dT)$ is omitted [85].

Rearranging terms, the temperature change dT after an evaporation step can be calculated as

$$dT = \frac{dU - \xi T_i k_B dN}{\xi N_i k_B}. \quad (7.5)$$

To determine the change in temperature dT after an evaporation step, both the change in internal energy dU and atom number dN must be determined.

We start by simulating the cloud of interacting bosons using the Bose statistics in Eq. (2.5), incorporating the measured harmonic trap frequencies of $(\omega_x, \omega_y, \omega_z) = 2\pi \times (102, 144, 232)$ Hz within an LDA framework, as described in Section 2.4. The temperature is set to $T = 260$ nK, derived from a fit of the virial expansion of the equation of state to the measured in-situ density profile (see Section 4.2). Interactions are incorporated using a self-consistent mean-field approach, adding the mean-field energy of the composite bosons to the phase-space dependent energy $E(\mathbf{r}, \mathbf{p})$. This results in a term of the total energy in the phase-space,

which reads

$$E(\mathbf{r}, \mathbf{p}) = p^2/(2m) + V(\mathbf{r}) + gn_\sigma(\mathbf{r}), \quad (7.6)$$

with $m = 2m_{\text{Li}}$ being the mass of the composite bosons, $V(\mathbf{r})$ being the harmonic potential (twice of what a single fermions experiences in Eq. (3.7)), $g = 2\pi\hbar^2 a/\mu_{\text{T}}$ [8] (refer to Eq. (2.23)) being the mean-field coupling constant with $\mu_{\text{T}} = m_{\text{Li}}$ (dimers) and n_σ being the density of composite bosons.

In order to determine the density distribution $n_\sigma(\mathbf{r})$, we integrate the Bose statistics in Eq. (2.5) over the momentum space $d^3\mathbf{p}$, as described in Eq. (2.31). Here, the energy term in Eq. (7.6) is considered in the LDA approach. This step is iterated self-consistently to incorporate the density dependent mean-field term until the density profile n_σ converges within an uncertainty of 1%. We can use the resulting density profile to recover the atom number N by integrating over the real space $d\mathbf{r}$.

The internal energy is determined by integrating the Bose statistics in Eq. (2.5) over the full phase space $d^3\mathbf{r}d^3\mathbf{p}$. By using the same integrand as for the calculation of the initial density profile earlier, the initial internal energy U_i is calculated. With repulsive interactions introducing slight deviations from a perfect harmonic confinement at the trap centre, we adjust ξ in Eq. (7.3) to fit the altered trap, finding $\xi = 2.81$, close to the harmonic trap value of $\xi = 3$. We use this value for calculations of the internal energy after removing atoms as well.

To simulate cooling, we modify the Bose statistics such that the probability of finding a boson for $r > r_{\text{cut}}$ drops to zero. This is expressed by

$$f_{\text{Bose}}(E(\mathbf{r}, \mathbf{p}), T) = \begin{cases} f_{\text{Bose}}(E(\mathbf{r}, \mathbf{p}), T) & \text{if } r < r_{\text{cut}} \\ 0 & \text{if } r \geq r_{\text{cut}}, \end{cases} \quad (7.7)$$

where $f_{\text{Bose}}(E(\mathbf{r}, \mathbf{p}), T)$ represents the Bose statistics in Eq. (2.5).

Finally, we calculate the final atom number N_f and internal energy U_f using the modified Bose statistics.[†] These values are then used in Eq. (7.5) to find the corresponding change in temperature dT .

We will apply this simulation in Section 7.3.1 to compare the observed magnitude of cooling with the predicted magnitude by this model.

7.3 Experimental implementation

In the previous section, we have discussed the concept and principles behind the new cooling technique. This section covers the implementation of this method and examines the impact of modulation parameters, specifically amplitude and frequency, on the observed cooling performance.

The cooling technique operates by dissociating composite dimers on the BEC side of the crossover and is implemented similarly to the spectroscopic measurements of the Higgs mode, discussed in Section 6.3.1. Dimers exist on the BEC side of BEC-BCS crossover, and have a bound state energy of E_B , as given by Eq. (2.11). Here, the magnetic field B determines the

[†]We still use the unaltered density distribution n_σ to incorporate the mean-field term.

position in BEC-BCS crossover and; consequently, defines the bound state energy E_B .

Dimer association is achieved by modulating the magnetic field using the custom-built small magnetic field coil introduced in Section 4.5.2. The resulting profile of the magnetic field is illustrated in Fig. 7.5.

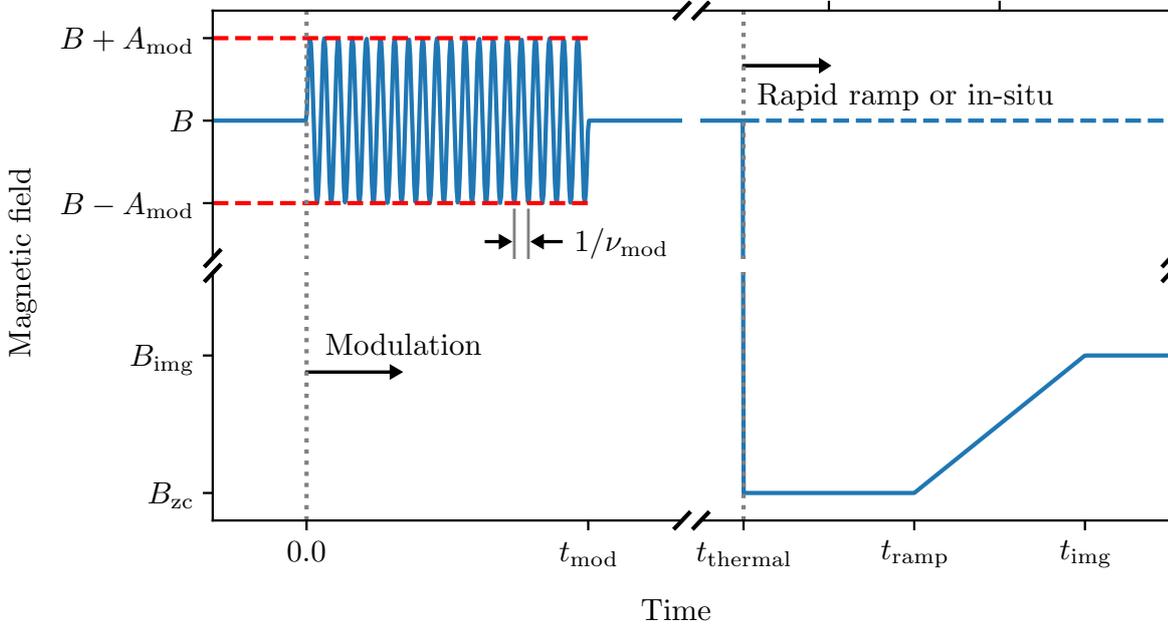


Figure 7.5: Experimental protocol for cooling via magnetic field modulation. This figure shows the magnetic field profile used to dissociate composite dimers through magnetic field modulation. The sample is first prepared at an offset field of B . At $t = 0$, a sinusoidal modulation with amplitude A_{mod} and frequency ν_{mod} is superimposed to the offset field for a duration of t_{mod} , using the small magnetic field coil introduced in Section 4.5. The modulation time t_{mod} is adjusted to contain an integer number $n_{\text{mod}} = 2\pi t_{\text{mod}} \nu_{\text{mod}}$ of oscillations. After modulation, a thermalisation period of $t_{\text{thermal}} - t_{\text{mod}}$ is added. To detect cooling, the sample is imaged either in-situ at t_{thermal} to measure the sample's temperature, or after additional expansion using the RR technique at t_{img} , to measure the condensate fraction CF , as detailed in Section 4.3. The rapid ramp includes two timescales: the free expansion time $t_{\text{ramp}} - t_{\text{thermal}}$ at the zero-crossing B_{zc} of the Feshbach resonance and the ramp time $t_{\text{img}} - t_{\text{ramp}}$ to the final imaging magnetic field B_{img} .

Initially, the sample is prepared at a magnetic field B , setting its position on the BEC side of the BEC-BCS crossover, and thus defining the bound state energy E_B . After preparing a $|13\rangle$ mixture as described in Section 3.2.6, cooling is achieved by superimposing the modulated magnetic field. The modulation parameters include an amplitude A_{mod} , frequency ν_{mod} and modulation duration t_{mod} . To prevent magnetic field jumps, the number of oscillations is ensured to stay integer, which is achieved by adjusting the modulation time to ensure $n_{\text{mod}} = 2\pi t_{\text{mod}} \nu_{\text{mod}}$.

Typically, after a modulation period of $t_{\text{mod}} \sim 300$ ms, a thermalisation time of $t_{\text{thermal}} - t_{\text{mod}} = 200$ ms is added to allow full thermalisation of the sample. We observe that the amount of dissociated atoms depends only on the combined expression $A_{\text{mod}}^2 t_{\text{mod}}$, consistent with [218], and $A_{\text{mod}}^2 t_{\text{mod}}$ also uniquely defines the extent of cooling. However, to avoid

introducing additional timescales, we only vary A_{mod} while keeping t_{mod} constant, unless stated otherwise. Note that the condensate has a $1/e$ decay time of approximately ~ 1 s. Later, in Section 7.5, we scale $t_{\text{mod}} \propto 1/\sqrt{A_{\text{mod}}}$ to probe the achievable cooling timescale with minimal modulation duration.

Probing the cooling technique

To quantify the cooling effect, we probe the cooled gas both in-situ and after the rapid ramp. The general detection procedure is detailed in Chapter 4, but a brief overview is provided here as well. To infer the 3D density of the trapped sample, the in-situ image is processed via the inverse Abel transformation, refer to Section 4.2.1. This processed 3D density provides the central density n_σ and the temperature T , determined by fitting the 3D density profile to the virial expansion of the equation of state (see Section 4.2). However, this procedure requires extensive data and is therefore supplemented by a faster measurement of the phase-space density via the condensate fraction CF , obtained from images taken after the rapid ramp (refer to Section 4.3).

While in-situ images are captured immediately after the thermalisation period, the rapid ramp incorporates an additional expansion time, as depicted in Fig. 7.5.

7.3.1 Observed cooling

First, we use the measurement of the condensate fraction CF as a proxy for the phase-space density to assess the influence of the modulation parameters, namely the amplitude A_{mod} and frequency ν_{mod} , on the cooling efficiency. To this end, we prepare a sample at a magnetic field of $B = 640.31(6)$ G, which corresponds to a bound state energy of $E_B/h = 104.5(5)$ kHz. In order to observe cooling, the magnetic field is modulated with frequencies close to this bound state energy. The observed change $CF - CF_0$ in condensate fraction from an initial condensate fraction CF_0 to a condensate fraction of CF is shown in Fig. 7.6.

The left plot of Fig. 7.6 shows the change in condensate fraction $CF - CF_0$ as a function of modulation amplitude A_{mod} for various modulation frequencies ν_{mod} . We observe different regimes of cooling and heating depending on the combination of amplitude A_{mod} and frequency ν_{mod} . For frequencies below the bound state energy $\nu_{\text{mod}} \leq E_B/h = 104.5(5)$ kHz, there is no discernible increase in condensate fraction $CF - CF_0$, and instead, the condensate fraction decreases at larger modulation amplitudes. Since dimers can not be broken at these frequencies, this depletion is attributed to thermal excitations from sound modes [48].

For frequencies above the bound state energy $\nu_{\text{mod}} > E_B/h = 104.5(5)$ kHz, we observe an increase in condensate fraction $CF - CF_0$ at lower modulation amplitudes A_{mod} , indicating an efficient cooling step. For much higher modulation amplitudes, however, the cooling decreases and eventually leads to heating. As shown later in Fig. 7.7, higher modulation amplitudes cause more atoms to be removed from the trap. Thus, the modulation amplitude A_{mod} of maximum cooling can be understood as the point at which we start to remove particles with an energy approaching the average energy per particle, effectively making evaporation less efficient.

In order to estimate the frequency ν_{mod} range for efficient cooling, we track the maximum change in condensate fraction $CF - CF_0$ across tightly sampled modulation frequencies, as

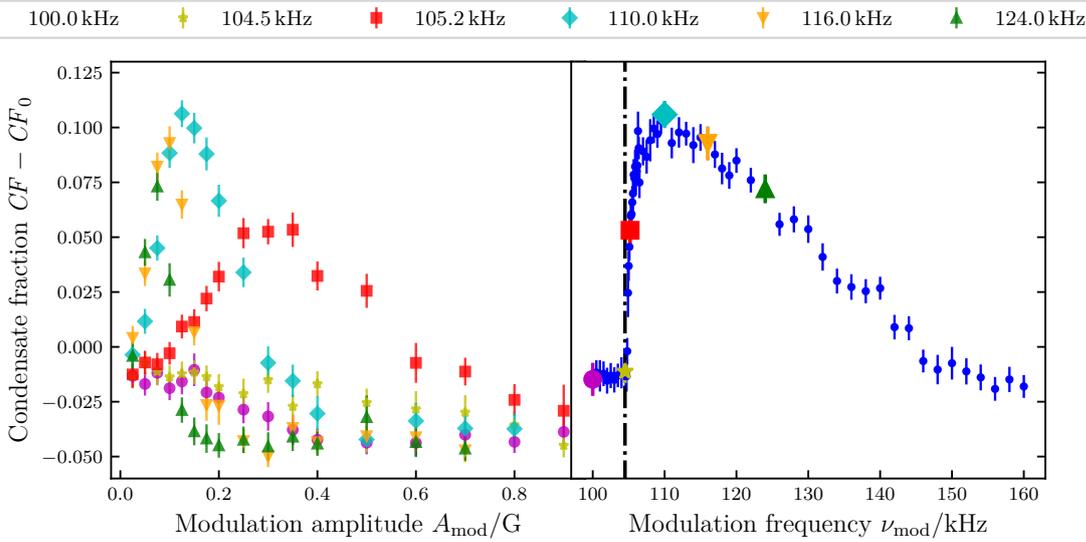


Figure 7.6: Condensate fraction after the magnetic field modulation. The figure presents the cooling effect as a change in condensate fraction $CF - CF_0$ relative to the initial condensate fraction CF_0 for an exemplary sample with bound state energy $E_B/h = 104.5(5)$ kHz (vertical dashed dotted). With a constant modulation time t_{mod} , the modulation has two free parameters: amplitude A_{mod} and frequency ν_{mod} . **Left:** Change in condensate fraction $CF - CF_0$ as a function of modulation amplitude A_{mod} for six exemplary modulation frequencies ν_{mod} . **Right:** Maximum change in condensate fraction for each modulation frequency, with the six frequencies from the **left** plot highlighted by corresponding markers and colours. Errorbars denote the standard error.

shown in the right plot of Fig. 7.6. Efficient cooling is possible for a wide range of frequencies, but becomes less efficient and eventually leads to heating at much higher frequencies. This can be attributed to the finite dissociation threshold of the coldest dimers in the centre of the trap, as given in Eq. (7.2). At very high frequencies, primarily the coldest, condensed dimers are broken, inhibiting cooling by dissociation. In other words, the dissociation method loses its selectivity.

Temperature reduction

We also evaluate the cooling technique by measuring the sample's temperature T . For this purpose, we prepare an equivalent sample at a magnetic field of $B = 640.31(6)$ G with a bound state energy of $E_B/h = 104.5(5)$ kHz. Again, we vary both the modulation amplitude A_{mod} and frequency ν_{mod} and measure the temperature T of the sample post-thermalisation. The measured temperatures are displayed in the left plot of Fig. 7.7. For a more direct comparison to the theoretical step model in Section 7.2.1, we replace the modulation amplitude A_{mod} on the horizontal axis with the fractional atom loss $-\Delta N/N_i$ with respect to the initial atom number N_i .[‡]

Similar to Fig. 7.6, we observe regimes of cooling and heating depending on the combination

[‡]Here ΔN is used instead of dN in Section 7.2.1 to acknowledge the non-infinitesimal loss of atoms.

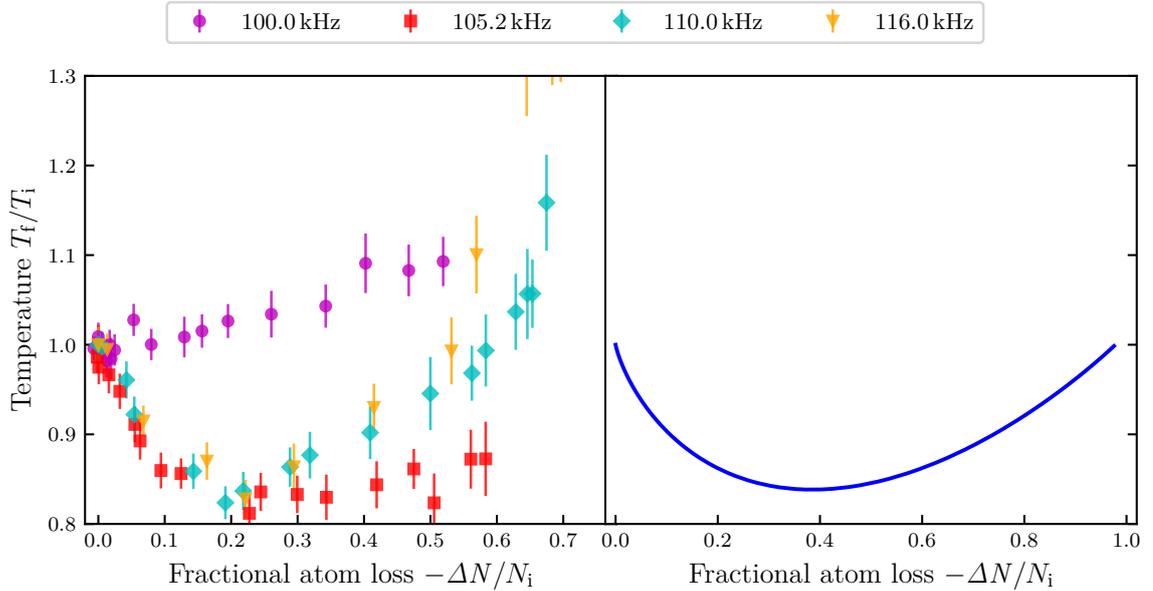


Figure 7.7: Temperature reduction through cooling. Analogous to Fig. 7.6, an exemplary sample with bound state energy of $E_B/h = 104.5(5)$ kHz is cooled via a magnetic field modulation from an initial temperature T_i . **Left:** For four exemplary modulation frequencies ν_{mod} , the temperature T_f is measured as a function of modulation amplitude A_{mod} , yielding a fractional atom loss of $-\Delta N/N_i$ from an initial atom number of N_i . Again, an onset at the bound state energy E_B is visible. **Right:** The measured temperature drop is compared with the simulated temperature reduction in an interacting Bose gas based on the step model of evaporative cooling, refer to Section 7.2.1. Here, a number $-\Delta N$ of atoms is removed from trap edge towards the centre. **Left:** Errorbars denote the standard error from the thermometry fit (refer to Section 4.2).

of frequency ν_{mod} and amplitude A_{mod} , i.e. the fraction atom loss $-\Delta N/N_i$. For frequencies below the bound state energy $\nu_{\text{mod}} < E_B/h = 104.5(5)$ kHz, no cooling is discernible, and instead, the sample heats up with increasing atom loss. However, for frequencies above the bound state energy $\nu_{\text{mod}} > E_B/h = 104.5(5)$ kHz, a small atom loss yields a reduced temperature of the sample, confirming the cooling effect. As discussed in the context of the condensate fraction CF earlier, a much larger atom loss leads to heating.

The theoretical step model presented in Section 7.2.1 is also shown in Fig. 7.7 for comparison. Consistent with our measurements, the model predicts a temperature reduction of approximately $\sim 15\%$ for a fractional atom loss around $\sim 30\%$. This strong agreement validates our description of the cooling mechanism as given in Section 7.2.

However, the model in Fig. 7.7 does not capture heating, which is observed for higher frequencies in the experimental data. This discrepancy can be attributed to two factors not included in the model. First, the model assumes that all atoms above the dissociation threshold are immediately expelled from the trap. In practice, however, dissociated fragments with excess kinetic may remain in the trap and heat the sample. Second, at much higher frequencies ν_{mod} and consequently greater atoms losses, more dimers from the condensate's zero-energy state are dissociated and potentially removed from the trap. This leads to rapid

decay of the condensate and heating of the sample.

7.3.2 Onset of dissociation and cooling

We have clarified the dependence of the modulation parameters – amplitude A_{mod} and frequency ν_{mod} – on the efficient regimes of cooling. Yet, the question remains: does the onset frequency of efficient cooling consistently match the bound state energy, or does a small, un-captured deviation persist beyond the model presented in Section 7.2? To address this, we probe the cooling technique by measuring the number of condensed atoms $N_{\text{condensed}}$ across various magnetic fields on the BEC side of the crossover, where a dimer bound state is present.

The measurement and analysis are shown in Fig. 7.8. In the upper row we display the increase in condensed atoms $N_{\text{condensed}}$ as a function of modulation frequency ν_{mod} for three exemplary magnetic fields B using a sufficient magnetic field modulation amplitude A_{mod} . To automate detection of the cooling onset, we calculate the point-to-point gradient, identifying the onset of cooling at the frequency where the gradient reaches a positive maximum.

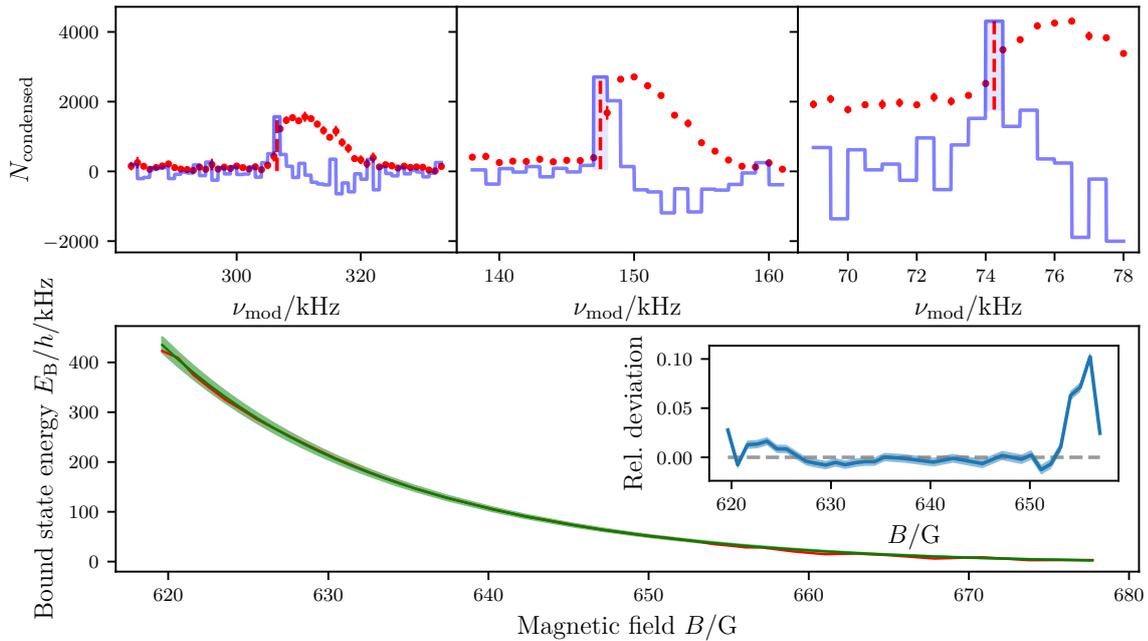


Figure 7.8: Onset of the cooling feature throughout the dimer regime of the BEC-BCS crossover. The figure highlights the measured onset of cooling via dimer dissociation, coinciding with the bound state energy E_B . **Upper row:** Cooling is identified by an increase in the number of condensed atoms $N_{\text{condensed}}$ (red points). For three exemplary samples prepared at magnetic fields of (left to right) 624.55(6) G, 635.38(6) G and 645.22(6) G, the onset (red dashed line) is identified by the maximum in the first derivative (blue line) of the data. **Bottom:** For each magnetic field, the measured onset (red line) is compared to the calculated bound state energy E_B at that field, calculated via the correction in Eq. (2.11). **Inset:** The relative deviation between the measured onset and calculated bound state energy is shown for quantitative comparison. Errorbars denote the standard error.

This analysis is applied to samples at various magnetic fields, yielding the magnetic field resolved onset frequency shown in the lower panel of Fig. 7.8. Alongside, we display the calculated bound state energy E_B , calculated with Eq. (2.11). We find excellent agreement between the detected onset frequency and the bound state energy.

For the largest magnetic fields – yielding the smallest bound state energies E_B – the relative deviation increases dramatically, attributed to the diminishing signal-to-noise ratio of the cooling peak, as the magnetic field approaches unitarity. The decline in cooling efficiency near unitarity will be further detailed in Section 7.4.3.

At the lowest magnetic fields, with the largest bound state energies E_B , the detected onset also shows stronger deviations from the calculated bound state energy. This can not be explained by a diminishing signal-to-noise ratio, as before, but by the limitations of the amplifier circuit of the small magnetic field coil, as discussed in Section 4.5.2.

Thus, we can conclude that, within our precision, the onset of the cooling technique coincides with the calculated bound state energy E_B , providing further validation of the mechanism behind the cooling technique described in Section 7.2.

7.3.3 Frequency of maximum cooling efficiency

Till now, we have examined the influence of the modulation amplitude A_{mod} and frequency ν_{mod} on our proposed cooling technique. We have observed that cooling occurs for modulation frequencies ν_{mod} slightly larger than the bound state energy of a sample $\nu_{\text{mod}} \geq E_B/h$. This aligns with the proposed cooling mechanism in Fig. 7.5, where the density dependent dimer dissociation threshold (see Eq. (7.2)) selectively removes high-energy atoms from the trap's edge.

The dimer dissociation threshold is highest at the centre, where the dimer density is greatest (see Eq. (7.1)), and predominantly zero-energy dimers from the condensate are found here. If the modulation frequency ν_{mod} is greater than the dimer dissociation threshold at the trap centre, many dimers from the condensate are dissociated and removed from the trap if their excess kinetic energy suffices. Dimers from the condensate have the lowest energy, and hence a removal would result in strong heating of the sample. Thus, the frequency range at which we expect efficient cooling should be limited to a frequency range similar to the dissociation threshold at the centre of the trap. Hence, the range of efficient cooling should depend on the density in the trap centre n_σ . To investigate, we tune the easily accessible parameter, the centre density n_σ , and probe the modulation frequency ν_{mod} at which we experience the most efficient cooling.

Measurements of the condensate fraction CF in Fig. 7.6 already reveal a modulation frequency ν_{mod} for maximum cooling efficiency. Here, we repeat the measurement in Fig. 7.6 with samples of varying initial density n_σ . We set the magnetic field to a value of $B = 627.50(6)$ G – the field furthest on the BEC side that still allows for modulation frequencies to be above the bound state energy of $E_B/h = 251.3(11)$ kHz. This maximises the relevance of dimer physics, ensuring that the dissociation threshold relation from Eq. (7.1) holds. Note that the interaction parameter for samples at this field is roughly $1/(k_F a) \sim 3$.

To adjust (reduce) the initial centre density n_σ , we vary (reduce) the initial preparation time in the optical dipole trap, refer to Section 3.2.6. Then, we perform the same measurement as outlined in Section 7.3.1, varying the modulation frequency ν_{mod} and amplitude A_{mod} . For

each modulation frequency ν_{mod} , we take the maximum value of the condensate fraction CF , which occurs at a finite modulation amplitude A_{mod} . We display the maximum condensate fractions CF for different initial centre densities n_σ as a function of modulation frequency ν_{mod} in Fig. 7.9.

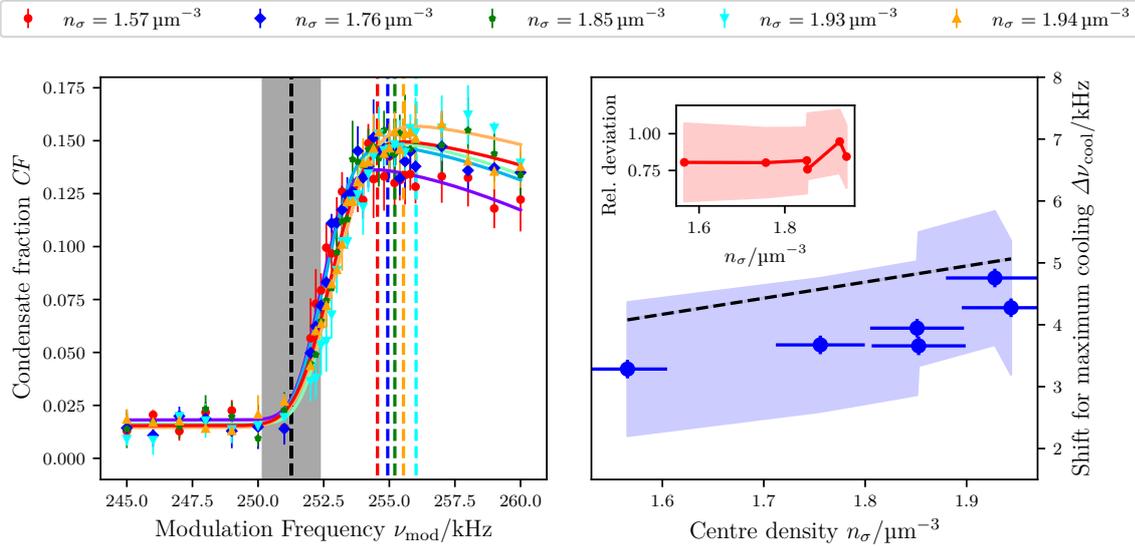


Figure 7.9: Density dependence of the frequency necessary for maximum cooling. As suggested by Fig. 7.4 and Eq. (7.1), the dissociation threshold frequency of a density bin n_σ is linearly shifted by density. To investigate this, we determine the frequency for maximum cooling efficiency in terms of the condensate fraction CF for samples as far on the BEC side as possible, with a bound state energy of $E_B/h = 251.3(11)$ kHz (vertical dashed band). By adjusting the initial preparation in the dipole trap (refer to Section 3.2.6), we probe the cooling technique for different initial centre densities n_σ (symbols and colours). **Left:** Higher centre densities shift the position of maximum CF to higher modulation frequencies $\nu_{\text{mod}} = E_B + \Delta\nu_{\text{cool}}$, quantified by fitting a deterministic skewed Gaussian model (defined in Appendix B). **Right:** We plot the determined frequency of maximum cooling $\Delta\nu_{\text{cool}}$ as a function of the centre density n_σ . Here, the statistical variation due to the fitted model is highlighted by the blue shaded band. For comparison, the dissociation shift (see Eq. (7.1)) corresponding to the centre density is shown as a black dashed line. **Inset:** For all probed initial centre densities, the relative deviation between the frequency $\Delta\nu_{\text{cool}}$ of maximum cooling and the calculated dissociation shift (black dashed line in **right**) is shown. Errorbars denote the standard error.

Across all densities n_σ , the onset of cooling begins at the bound state energy E_B , as anticipated. At higher modulation frequencies ν_{mod} , CF increases, featuring a peak at a finite modulation frequency of $\nu_{\text{mod}} = E_B/h + \Delta\nu_{\text{cool}}$. In order to extract the position of maximum cooling $\Delta\nu_{\text{cool}}$, we fit a deterministic skewed Gaussian model to the data, defined in Appendix B. We fix the skewness parameter of the model to make the fitting routine more robust.

After fitting, we extract the maximum position of the skewed Gaussian model, yielding the position of maximum cooling $\Delta\nu_{\text{cool}}$. We display the determined positions $\Delta\nu_{\text{cool}}$ for all measured densities n_σ in the right panel of Fig. 7.9. Due to the uncertainty of the fitting routine, the position of maximum cooling $\Delta\nu_{\text{cool}}$ has significant uncertainty, highlighted by

the blue shaded band. Alongside the extracted data, we show the calculated (see Eq. (7.1)) dissociation threshold frequency at the trap centre as a reference. We find that the frequency $\Delta\nu_{\text{cool}}$ at which we observe maximum cooling depends linearly on the density n_σ at the centre of the trap. We expect the same proportionality for the dissociation threshold frequency in Eq. (7.1).

Moreover, the maximum cooling frequency $\Delta\nu_{\text{cool}}$ consistently occurs at around $\sim 80\%$ of the calculated dissociation threshold frequency at the trap centre, independent of density n_σ . This confirms the linear dependence of the position of maximum cooling $\Delta\nu_{\text{cool}}$ with centre density n_σ .

Notably, while the frequency of maximum cooling $\Delta\nu_{\text{cool}}$ linearly depends on the centre density n_σ , the line shape for efficient cooling is broad, which allows efficient cooling over a wide range of frequencies. Although all dimers with a dissociation threshold below ν_{mod} can indeed dissociate, only those with sufficient excess kinetic energy are immediately removed from the sample, directly impacting temperature through evaporation. Dimers with a threshold frequency at or only slightly above the modulation frequency ν_{mod} can be dissociated, but remain in the trap, interacting with the remaining sample. This effect is not captured by our developed step model in Section 7.2.1, and the subsequent dynamics are much more complicated. Hence, $\Delta\nu_{\text{cool}}$ primarily reflects a frequency range of efficient cooling, while the broader line shape reflects more complex behaviour.

7.4 Efficiency of cooling

So far, we have investigated the working principles of the cooling technique and have demonstrated cooling for a wide range of modulation parameters – frequency and amplitude. To gain a more quantitative understanding, we now examine the efficiency γ of this cooling technique. The efficiency γ is a key metric for any evaporative cooling technique [85], as it represents the expected increase in phase-space density $\tilde{\rho}$ during the evaporative loss of $-\Delta N = N_f - N_i$ atoms. Since maximising the phase-space density $\tilde{\rho}$ is the goal of any evaporative cooling approach, a high efficiency γ is desirable. Measuring the cooling efficiency also allows for a direct comparison to alternative cooling methods.

The efficiency of an evaporative cooling step is defined as [85]

$$\gamma = -\frac{\log(\tilde{\rho}_f/\tilde{\rho}_i)}{\log(N_f/N_i)}, \quad (7.8)$$

with a change of the phase-space density from $\tilde{\rho}_i \rightarrow \tilde{\rho}_f$ and a change of atoms from $N_i \rightarrow N_f$. An efficiency of $\gamma > 1$ is considered effective cooling, while a negative cooling efficiency $\gamma < 0$ would indicate heating. Here, the phase-space density $\tilde{\rho}$ is defined as

$$\tilde{\rho} = n_\sigma \lambda_{\text{dB}}^3, \quad (7.9)$$

with n_σ as the sample density (per spin state) and λ_{dB} as the thermal de-Broglie wavelength.

We determine the efficiency by probing the sample in-situ after cooling, as described in Section 7.3. From in-situ images, we obtain the density profile $n_\sigma(r)$ as well as the sample temperature T , outlined briefly in Section 7.3 and in more detail in Section 4.2. From the

density profile $n_\sigma(r)$ we infer the density at the trap centre $n_\sigma := n_\sigma(r=0)$, which allows us to determine the phase-space density $\tilde{\rho}$. We also use the density at the centre of the trap n_σ to gauge the Fermi temperature T_F for a relative temperature scale (see Section 2.4).

The measured temperature T/T_F and phase-space density $\tilde{\rho}$ are shown for an exemplary magnetic field of $B = 640.31(6)$ G in Fig. 7.10. This magnetic field, with a corresponding bound state energy of $E_B/h = 104.5(5)$ kHz, is the same as in Fig. 7.6 and Fig. 7.7. To measure the maximum efficiency γ possible for this field, we set the modulation frequency to $\nu_{\text{mod}} = 110$ kHz, where the maximum condensate fraction is observed in Fig. 7.6. In that figure, we also note that the maximum cooling efficiency for each modulation frequency ν_{mod} occurs at a specific modulation amplitude A_{mod} , corresponding to a specific fractional atom loss $-\Delta N/N_i$. Thus, we measure the temperature T/T_F and phase-space density $\tilde{\rho}$ for various modulation amplitudes A_{mod} . For a clearer analysis, we replace the modulation amplitude A_{mod} on the horizontal axis with the fraction atom loss $-\Delta N/N_i$, as done before in Fig. 7.7.

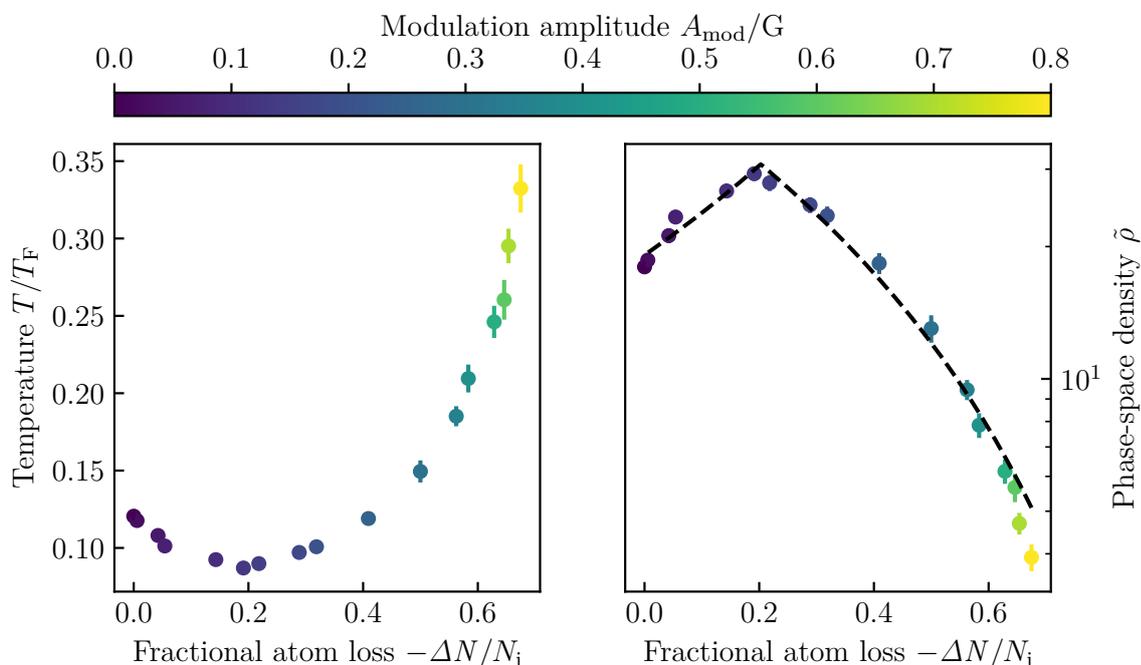


Figure 7.10: Determination of the cooling efficiency. To quantify the cooling method, we determine the efficiency γ of the phase-space density increase $\tilde{\rho}$ relative to the fraction atom loss $-\Delta N/N_i$. For this purpose, an exemplary sample with a bound state energy of $E_B/h = 104.5(5)$ kHz is cooled through magnetic field modulation at a frequency of $\nu_{\text{mod}} = 110$ kHz, corresponding to the maximum in Fig. 7.6. The amplitude A_{mod} of the modulation is changed between measurements, producing a fractional atom loss of $-\Delta N/N_i$. **Left:** As seen in Fig. 7.7, the temperature drops for a moderate atom loss, indicating cooling. **Right:** Concurrently, the phase-space density $\tilde{\rho}$ increases. We fit the model from Eq. (7.8) in two piecewise segments to the data to extract the cooling efficiency γ from the initial positive slope segment in the logarithmic plot. Errorbars denote the standard error.

For small fractional atom losses $-\Delta N/N_i$, we observe a temperature decrease and a phase-space density $\tilde{\rho}$ increase. As previously noted and discussed in Fig. 7.6 and Fig. 7.7, we

experience less efficient cooling and ultimately heating for a much stronger atom loss. Efficient cooling is observed until where the maximum phase-space density $\tilde{\rho}$ is reached, at a fractional atom loss of approximately $-\Delta N/N_i \sim 0.2$.

To determine the cooling efficiency γ , we fit Eq. (7.8) in a piece-wise fashion of two segments to the data. The first segment of positive slope describes efficient cooling, while the second segment of negative slope indicates heating with $\gamma < 0$. For this exemplary field, we deduce a cooling efficiency of $\gamma = 2.10(23)$.

7.4.1 Frequency range of efficient cooling

In the previous section, we have determined the cooling efficiency γ for an exemplary sample at a magnetic field of $B = 640.31(6)$ G with a corresponding bound state energy of $E_B/h = 104.5(5)$ kHz. To maximise the possible efficiency γ , we have employed a modulation frequency ν_{mod} at which we observe the maximum condensate fraction CF in Fig. 7.6. While this is a reasonable approach, a full spectrum of the cooling efficiency $\gamma(\nu_{\text{mod}})$, similar to Fig. 7.6, would provide more insight into the range of possible cooling efficiencies.

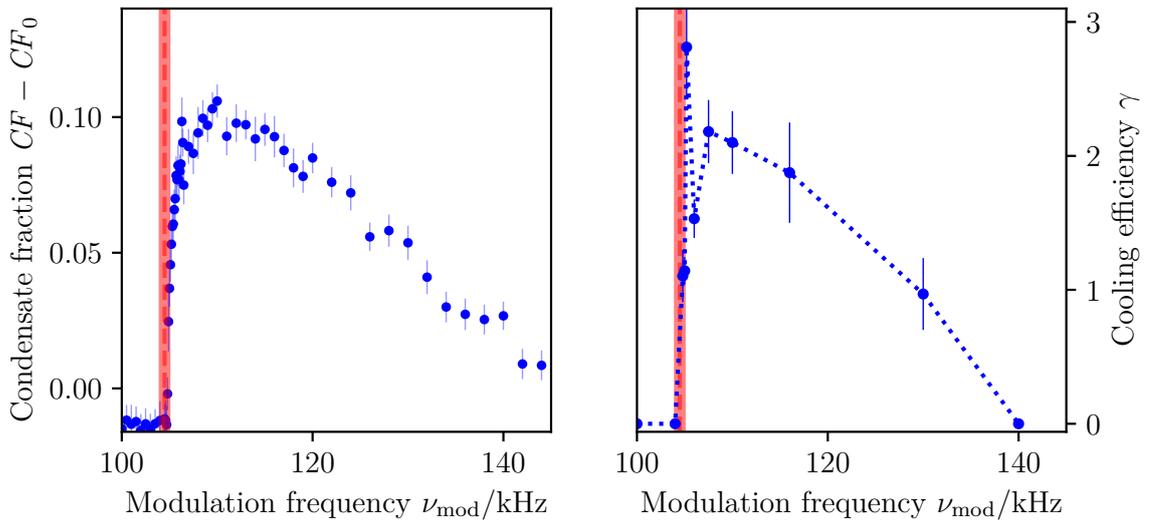


Figure 7.11: Linewidth of the increase of the condensate fraction and the cooling efficiency. To quantify the cooling technique, we measure both the condensate fraction CF (**left**) from RR images (see Section 4.3), and the temperature and density from in-situ images (see Section 4.2). The latter method allows extraction of the cooling efficiency γ (**right**). The figure shows both the increase in $CF - CF_0$ (**left**) as well as in γ (**right**) for a sample prepared with a bound state energy of 104.5(5) kHz (vertical red line) and cooled with a modulated magnetic field at a frequency of ν_{mod} . Errorbars denote the standard error.

To this end, we have measured the cooling efficiency γ as described in Section 7.4 but across several modulation frequencies ν_{mod} at the same exemplary magnetic field of $B = 640.31(6)$ G. This produces a spectrum of cooling efficiencies $\gamma(\nu_{\text{mod}})$ akin to the spectrum of the maximum change in condensate fraction $CF - CF_0$, as shown in Fig. 7.6. Both spectra of γ and $CF - CF_0$ are presented in Fig. 7.11. Note that only the positive cooling efficiency $\gamma > 0$ of the first

segment – with moderate atom loss – is shown. If no cooling is discernible – and the phase-space density graph in Fig. 7.10 can be described by a single segment – the cooling efficiency is set to $\gamma = 0$.

The spectrum confirms that, while in both γ and $CF - CF_0$ regimes of efficient and less efficient cooling are discernible, the frequency range of maximum efficient cooling is broad. Thus, desiring the best cooling efficiency does not require an extremely accurate determination of the best setting for the modulation. Also, the spectrum confirms that the maximum cooling efficiency γ occurs at the frequency at which we observe the maximum increase in condensate fraction. Therefore, we can stick with the simpler determination of the maximum efficiency by measuring the condensate fraction CF .

7.4.2 Cooling throughout the BEC regime of the BEC-BCS crossover

Using insights from previous measurements, we now investigate the cooling technique throughout the BEC regime of the BEC-BCS crossover by varying the magnetic field and, thereby, the bound state energy E_B . To this end, we measure the maximum change in condensate fraction $CF - CF_0$ as shown in the right panel of Fig. 7.6. As shown in Fig. 7.11, the modulation frequency ν_{mod} that yields the maximum increase in $CF - CF_0$ aligns with the highest achievable cooling efficiency γ . Consequently, we measure the cooling efficiency γ for various bound state energies E_B using the modulation frequency at which the gain in condensate fraction $CF - CF_0$ is maximal. The results for both the maximum gain in $CF - CF_0$ and the corresponding maximum cooling efficiency γ are shown in Fig. 7.12.

The measurement suggests that the cooling method is generally more effective for more deeply bound dimers with higher bound state energies E_B . We observe most efficient cooling for bound state energies close to and above $E_B/h \sim 200$ kHz. For these samples, we achieve a maximum gain in the condensate fraction of $CF - CF_0 \sim 0.15$, while the cooling efficiency reaches high values of $\gamma \sim 4$. This efficiency is exceptionally high and comparable to, or even surpasses, the largest reported evaporation efficiencies of $\gamma \leq 4$ in forced evaporation approaches in an optical dipole trap near a Feshbach resonance [85, 87, 91–99].

Conversely, as the bound state energy approaches zero $E_B \rightarrow 0$ near unitarity, cooling becomes significantly less efficient. For a bound state energy of $E_B/h = 73.5(5)$ kHz we achieve a cooling efficiency of only $\gamma = 1.08(16)$, barely exceeding unity. This situation occurs for an interaction parameter of $1/(k_F a) \sim 0.7$ and a magnetic field of $B = 645.22(6)$ G. As we move to fields with even lower bound state energies, the cooling efficiency falls below unity, rendering the cooling method inefficient.

We can attribute this reduction in efficiency to two factors, which align with our proposed cooling mechanism in Section 7.2. First, as E_B decreases, the gap in dissociation threshold (see Eq. (7.2)) between the high- and low-energy dimers becomes less pronounced, weakening the selective nature of the cooling method, and thus reducing its effectiveness. Second, as the dimer pair size increases approaching unitarity [17], the behaviour of the dimer sample increasingly resembles the dynamics of overlapping many-body interactions rather those of a weakly-interacting bosonic dimers gas. This shift affects the spectral line shape of dissociation [17] and, consequently, influences the cooling technique. A more precise determination of the magnetic field beyond which no cooling effect is observed is carried out in Section 7.4.3.

Note that these measurements have been performed in a gravity-tilted trap. Efficiency

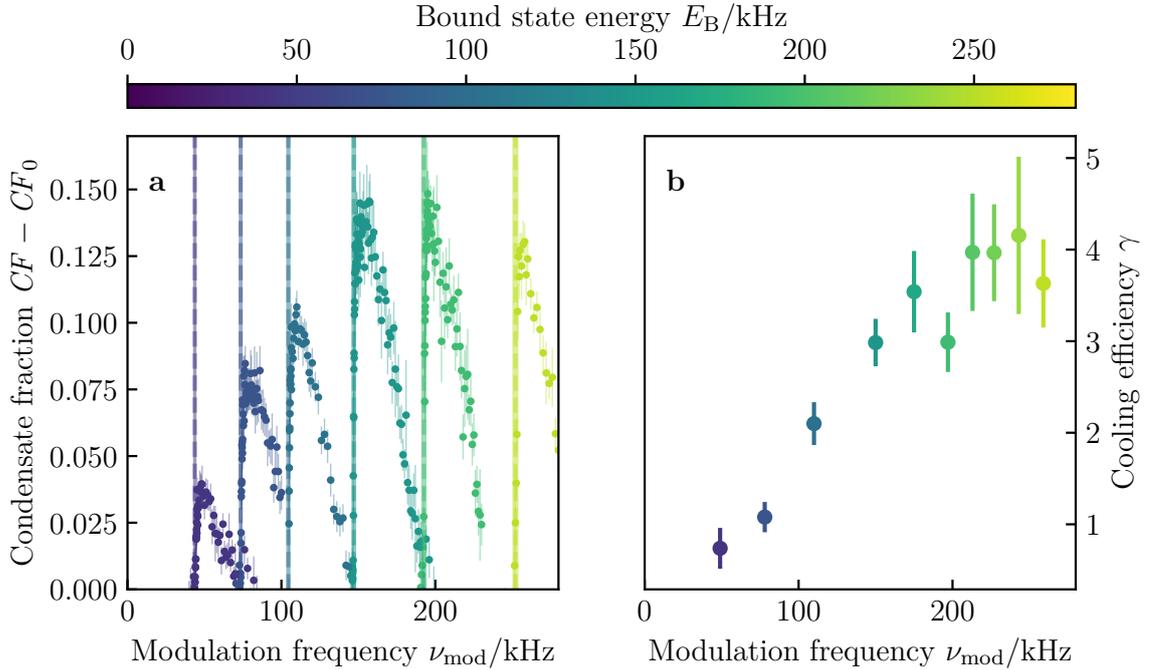


Figure 7.12: Cooling efficiency throughout the BEC side of the BEC-BCS crossover. The cooling technique operates via dimer dissociation and is thus limited to the bosonic side of the BEC-BCS crossover. Near unitarity, the cooling efficiency declines and eventually fails to cool the sample. The figure shows both the evolution of the condensate fraction increase $CF - CF_0$ (a) and the maximum cooling efficiency γ (b) for samples of different bound state energies E_B (colour). **a** Displays $CF - CF_0$ as determined in Fig. 7.6, with vertical lines indicating the respective bound state energy E_B . **b** For each E_B in **a**, the frequency of the maximum $CF - CF_0$ is employed to measure the maximum cooling efficiency γ , determined as shown in Fig. 7.10. Errorbars denote the standard error.

measurements in a gravity-compensated trap are presented in 7.4.4, where we observe comparable and even slightly higher efficiencies.

7.4.3 Minimum bound state supporting efficient cooling

As discussed in Section 7.4.2, the cooling efficiency γ decreases significantly and eventually falls below $\gamma = 1$ as the magnetic field approaches the unitarity limit. In the previous section, we already identified reasons for a reduced cooling efficiency, as the bound state energy $E_B \rightarrow 0$ approaches zero near unitarity. However, the sampling resolution of the efficiency measurement in Fig. 7.12 is insufficient to precisely determine the magnetic field B at which cooling becomes undetectable ($\gamma = 0$) and is replaced by heating only.

Here, we perform a measurement of higher resolution to identify the magnetic field B at which the cooling technique ceases to provide cooling. To this end, we measure the number of condensed atoms $N_{\text{condensed}}$ from the RR measurement and search for an increase in the number of condensed atoms $N_{\text{condensed}}$ for modulation frequencies ν_{mod} larger than the bound state energy E_B . The results are presented in Fig. 7.13.

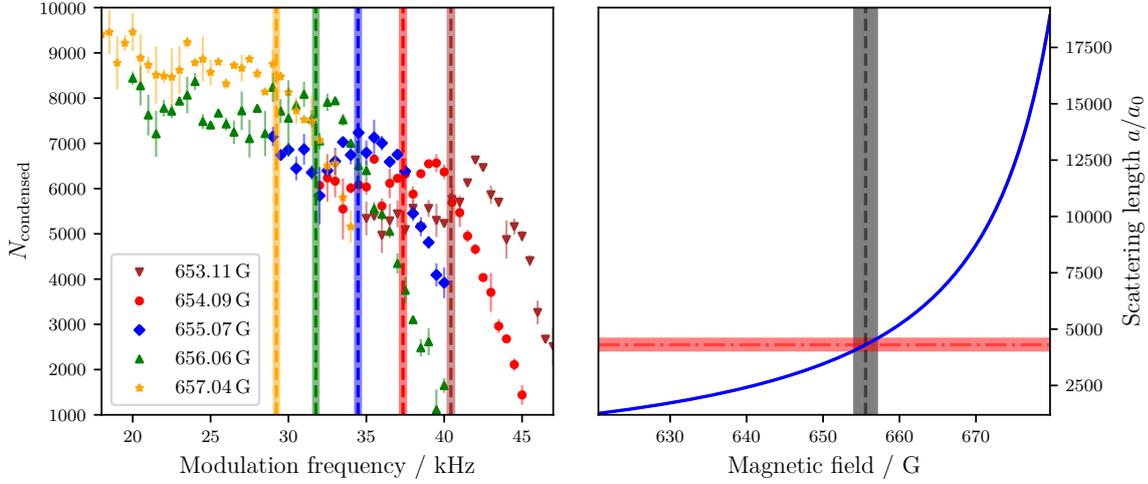


Figure 7.13: Vanishing of cooling when approaching unitarity. As indicated in Fig. 7.12, the cooling technique becomes less efficient towards unitarity and for very low bound state energies stops cooling the sample at all. In this figure we search for the highest magnetic field B or lowest bound state energy E_B , at which no cooling is discernible anymore. **Left:** To this end, we measure the amount of condensed atoms $N_{\text{condensed}}$ after the RR as a function of modulation frequency ν_{mod} and for fields that are sampled around the expected limiting highest field. Vertical lines highlight the corresponding bound state energies E_B . Between the smallest magnetic field (largest E_B) and the largest magnetic field (smallest E_B), we observe a trend from an increase in $N_{\text{Condensed}}$ to a featureless decay. **Right:** From this observation we pinpoint the onset of cooling between the measured fields (vertical grey area), corresponding to a range in the two-body scattering length a (horizontal red area). Errorbars are standard errors.

In the left panel of Fig. 7.13, we show the number of condensed atoms $N_{\text{condensed}}$ after a cooling step of modulation frequency ν_{mod} and a fixed modulation amplitude of $A_{\text{mod}} = 0.3$. For the smallest magnetic field B , or the largest bound state energy E_B , we still see a small remaining gain in the number of condensed atoms $N_{\text{condensed}}$. However, for the largest magnetic field B , or the smallest bound state energy E_B , no gain is observed. Between these magnetic fields, a smooth transition from cooling to heating occurs.

We estimate this transition to be between magnetic fields of $B \sim 654$ G and ~ 657 G, corresponding to a range in bound state energies from $E_B/h \sim 29$ kHz to ~ 37 kHz or a range in scattering lengths from $a/a_0 \sim 4051$ to ~ 4574 . Beyond this magnetic field, no cooling can be observed.

7.4.4 Efficiency in a gravity-compensated trap

In Section 7.4.2, we have measured the efficiency γ across various bound state energies E_B throughout the BEC side of the BEC-BCS crossover. This measurement has been performed in an optical dipole trap with a tilt due to gravity. Gravity is often beneficial for forced evaporative cooling in an optical dipole trap due to the reduced trap depth, which accelerates evaporation [85]. While gravity can enhance the efficiency of evaporation and the maximum achievable phase-space density $\tilde{\rho}$, it usually also decreases the total amount of trapped atoms

due to the lower trap depth, influencing any evaporative cooling technique significantly.

In order to investigate gravity's role in our cooling technique, we conduct an additional measurement of the cooling efficiency γ in a gravity-compensated trap. Here, gravity is countered by a magnetic gradient field. The measurement is performed similarly to that shown in Fig. 7.10, with the results for the gravity-compensated trap presented in Fig. 7.14.

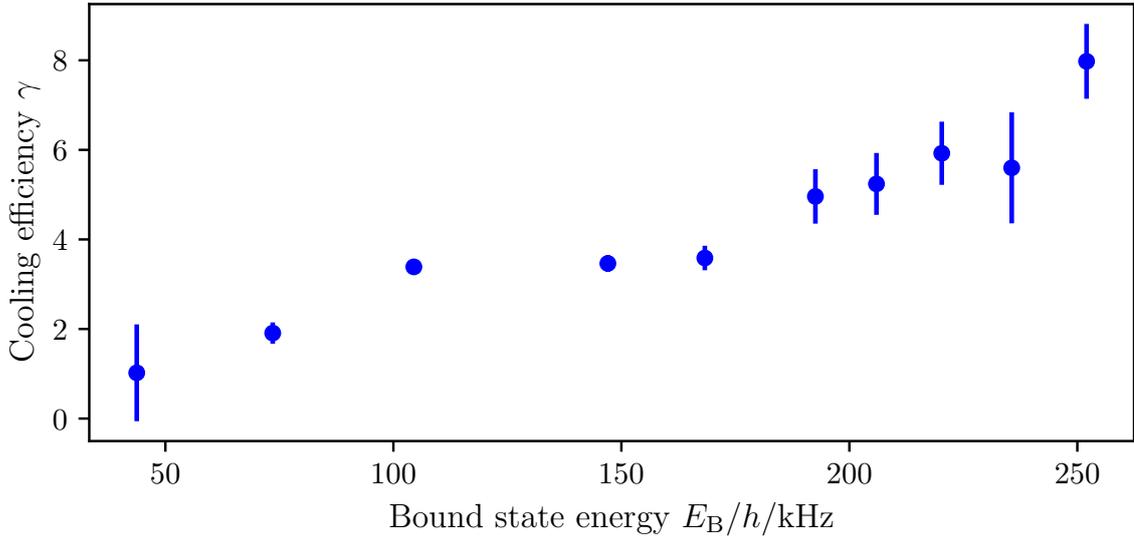


Figure 7.14: Cooling efficiency throughout the BEC side of the crossover for a gravity-compensated trap. In Fig. 7.12, we have shown the cooling efficiency γ throughout the bosonic side of the BEC-BCS crossover for a gravity-tilted trap. In this figure, we measure the cooling efficiency for the same frequencies as in Fig. 7.12, but for a gravity compensated trap, reaching comparable and even slightly higher efficiencies γ . Errorbars denote the standard error.

Qualitatively, the results align with those in the gravity-tilted trap, achieving the highest efficiencies γ for the highest bound state energies E_B . We observe efficiencies γ similar to those measured in the gravity-tilted trap in Fig. 7.10. More specifically, efficiencies range from $\gamma \sim 5$ to ~ 6 , slightly surpassing the maximum efficiency of $\gamma \sim 4$ in the gravity-tilted trap. Higher efficiencies are most likely an effect of the generally higher atom number in gravity-compensated traps, improving thermalisation and the efficiency of cooling. For the largest bound state energy, we measure a cooling efficiency close to $\gamma = 8$, which possibly presents a statistical outlier. These cooling efficiencies outperform the reported cooling efficiencies of up to $\gamma \leq 4$ [85, 87, 91–99] by a factor of up to 2.

While the cooling efficiencies γ are generally higher in the gravity-compensated trap, the absence of a tilt alters the trap geometry, leading to lower overall phase-space densities $\tilde{\rho}$. In a gravity-tilted trap, we observe a maximum gain in the condensate fraction of $CF - CF_0 \sim 0.1$ for a bound state energy of $E_B = 104.5(5)$ kHz in Fig. 7.12, resulting in a total condensate fraction of $CF \sim 0.16$. By contrast, in a gravity-compensated trap at the same bound state energy $E_B/h = 104.5(5)$ kHz, we measure a maximum change in condensate fraction of $CF - CF_0 \sim 0.04$, yielding a total condensate fraction of just $CF \sim 0.05$. Direct measurements of the phase-space density $\tilde{\rho}$ confirm these findings.

Given that the cooling technique is applied to a pre-cooled sample, the significantly higher achievable phase-space densities $\tilde{\rho}$ in a gravity-tilted trap may be more advantageous, despite a slightly reduced cooling efficiency.

7.5 Time scale of the cooling technique

So far, we have discussed the working principle of the cooling technique in Section 7.2, examined the dependency of the cooling efficiency on the parameters of the magnetic field modulation in Section 7.3 and determined the cooling efficiency for a broad range of bound state energies in Section 7.4. As outlined in Section 7.2, the cooling mechanism operates through the dissociation of high-energy dimers, potentially ejecting the dissociated fragments from the trap. We have observed the ejection of dissociated atoms in Fig. 7.4. The dissociation and subsequent separation of ejected fragments establishes the timescale of the cooling technique, followed by an additional thermalisation period to settle for the reduced temperature. Investigating this timescale allows for an optimised integration of this cooling technique in various systems, where time constraints may be strict due to limited sample lifetimes.

The cooling technique uses a modulation of the magnetic field with amplitude A_{mod} , frequency ν_{mod} and a time t_{mod} . To investigate the timescale on which the sample is cooling, the modulation time t_{mod} should be minimised to reduce smearing or overlap between t_{mod} and the time which it takes for the reduced temperature to settle. In Section 7.3, we have discussed that the amount of dissociated atoms only depends on the combined expression $A_{\text{mod}}^2 t_{\text{mod}}$, which relates amplitude A_{mod} and time t_{mod} , in line with earlier studies in [100, 218]. Thus, it is reasonable to assume that the cooling efficiency also depends on $A_{\text{mod}}^2 t_{\text{mod}}$.

To determine the timescale, we measure the time-resolved gain in condensate fraction $CF - CF_0$ for several modulation times $t_{\text{mod}} \propto 1/\sqrt{A_{\text{mod}}}$ at constant $A_{\text{mod}}^2 t_{\text{mod}}$ in Fig. 7.15. Here, we choose a value for $A_{\text{mod}}^2 t_{\text{mod}}$ at which we have observed maximum cooling before in Fig. 7.6. For each modulation time t_{mod} , we measure the time resolved gain in condensate fraction by varying the hold time t_{hold} in the dipole trap after cooling and before imaging. The hold time replaces the thermalisation time (refer to Fig. 7.5) – for $t_{\text{hold}} = 0$ an image is taken directly at the end of the modulation of the magnetic field. Due to the amplitude limit of the small magnetic field coil amplifier used for the modulation (refer to Section 4.5.2), the shortest possible modulation time is $t_{\text{mod}} = 9.3$ ms, which is significantly shorter than the typical 300 ms.

First, we observe that the gain in condensate fraction $CF - CF_0$ reaches the same level across all modulation times t_{mod} . This confirms that the optimal modulation time scales as $t_{\text{mod}} \propto 1/\sqrt{A_{\text{mod}}}$, as anticipated. This indicates that the cooling method can be applied much faster if a sufficiently large modulation amplitude is available.

Second, for the shortest possible modulation time of $t_{\text{mod}} = 9.3$ ms, we observe a time-resolved increase in condensate fraction $CF - CF_0$ after completing the magnetic field modulation. This dataset allows us to deduce the thermalisation timescale necessary to settle for the reduced temperature. To estimate this, we fit a deterministic model to the data, which captures the timescale of settling for the increased condensate fraction. We choose to fit the model in Eq. (6.6), which we have used earlier to describe the revival of the

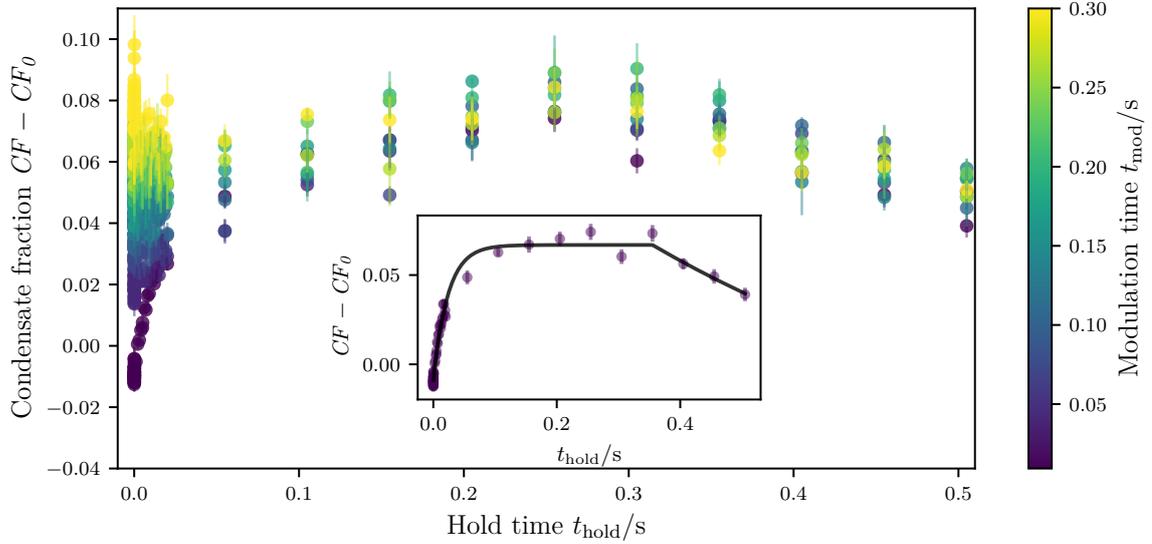


Figure 7.15: Time dynamics after modulation. The cooling technique works by dimer dissociation, subsequent ejection of high-energy atoms from the trap and re-thermalisation of the remains, and thus is connected to a finite timescale. In order to determine the timescale of the fastest possible cooling step, we employ the cooling technique for different modulation times t_{mod} of constant $A_{\text{mod}}^2 t_{\text{mod}}$ and measure the increase of the condensate fraction $CF - CF_0$ as a function of the hold time t_{hold} before the RR. **Inset** For the shortest possible modulation time of $t_{\text{mod}} = 9.3$ ms, we fit (solid line) a deterministic model from Eq. (6.6) to the data, revealing the timescale. Errorbars denote the standard error.

condensate fraction after a sudden quench of the magnetic field. From the fit, we find a $1/e$ thermalisation time of 23.8(11) ms. Hence, we can fully implement the cooling method in a time of less than 35 ms, incorporating $t_{\text{mod}} = 9.3$ ms of modulation time and 23.8(11) ms of thermalisation.

After an additional period of 355(29) ms, we begin to detect decay of the condensate. By fitting the thermal decay with and without modulation, we determine a $1/e$ thermal decay time of 0.92(1) s for the condensate in our setup.

7.6 Conclusion

We have introduced a novel cooling scheme for an interacting Fermi gas on the BEC side of the BEC-BCS crossover in an inhomogeneous optical dipole trap potential. This cooling technique selectively dissociates and ejects the highest energy dimers in the trap by modulating the magnetic field at frequencies above the corresponding dimer bound state energy. As the dissociation and ejection occur via additional magnetic field modulation, the trapping potential remains unaffected, simplifying the process of staying in the runaway regime of high cooling efficiencies.

The cooling technique relies on a sufficiently deep bound state, so that the dissociation threshold difference between the highest and lowest energy dimers allow targeted removal

of predominantly high energy dimers. Thus, the cooling technique performs best at large bound state energies on the BEC side of the crossover, where we have measured the largest cooling efficiencies γ . We have measured cooling efficiencies of up to $\gamma = 4$ in a gravity-tilted trap and even slightly larger efficiencies of up to $\gamma = 6$ in gravity-compensated trap.

Achieving optimal cooling efficiencies requires setting appropriate modulation parameters: frequency ν_{mod} , amplitude A_{mod} , and time t_{mod} . In Section 7.5, we have figured out that the latter two parameters can be adjusted interchangeably as $t_{\text{mod}} \propto 1/\sqrt{A_{\text{mod}}}$, with optimal cooling depending solely on the combined expression $A_{\text{mod}}^2 t_{\text{mod}}$ and frequency ν_{mod} . This enables rapid cooling with a modulation time as short as $t_{\text{mod}} = 9.3$ ms, followed by a $1/e$ thermalisation time of ~ 20 ms.

In this chapter, we have focused our analysis of the cooling technique on pre-cooled atoms in a $|13\rangle$ mixture. To explore the applicability of the new method, we have also applied the cooling technique to a $|12\rangle$ mixture with similar bound state energies E_{B} and observed comparable results as presented in Section 7.3. Additionally, we assessed the technique's effectiveness in an initially thermal gas with no condensate fraction. To this end, we have measured the temperature reduction similarly to Fig. 7.10 but for a single fixed modulation amplitude of $A_{\text{mod}} = 0.3$ G to reduce the data requirements. In order to start with an initially thermal gas, we have reduced the initial preparation time in the optical dipole trap (refer to Section 3.2.6). As shown in Fig. 7.16, cooling is also observed in an initially thermal dimer gas with no condensate, measured for a sample with a bound state energy of $104.5(5)$ G, matching the sample from Fig. 7.10.

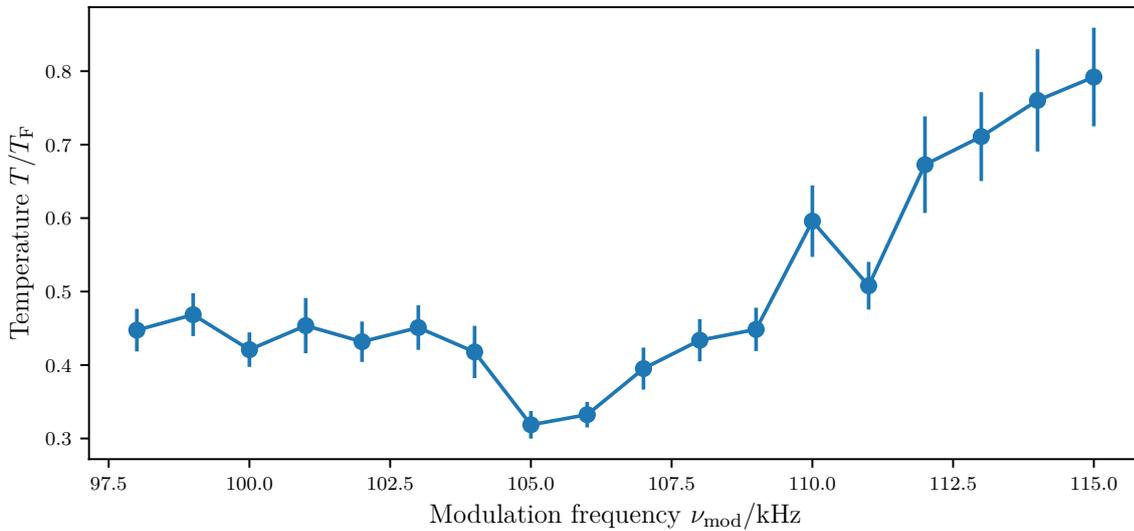


Figure 7.16: Cooling of an initially thermal gas. Considering the description of the cooling technique in Fig. 7.4, the requirement for the cooling technique to operate is a sufficiently deep two-body bound state. Thus, the cooling technique should work as long as the initial temperature supports this bound state. To investigate this, we measure cooling in terms of the temperature T/T_{F} as a function of the modulation frequency ν_{mod} and a constant modulation amplitude of $A_{\text{mod}} = 0.3$ G for a thermal sample with a bound state energy of $E_{\text{B}}/h = 104.5(5)$ kHz. Errorbars denote the standard error.

In conclusion, the method proves to be applicable in many systems, reaching high cooling

efficiencies throughout a wide region of the BEC regime. Since the implementation is not tied to many prerequisites, and the implementation of magnetic field modulation does not require huge alterations of existing systems [163], we expect the cooling method to be applicable in many systems.

Summary and Outlook

8.1 Summary

The key results of this thesis are presented in Chapter 5, Chapter 6, and Chapter 7. To conclude, we provide a concise summary of these chapters, highlighting the main findings and discussing the final outcomes.

Machine learning of the superfluid phase transition

In Chapter 5, we have investigated the superfluid phase transition in a fermionic spin mixture across the BEC-BCS crossover. The primary aim of this chapter was to determine a highly accurate measurement of the superfluid critical temperature within the BEC-BCS crossover. To achieve this, we have employed two distinct machine learning approaches to detect the phase transition directly from time-of-flight images.

The first method utilises a deep convolutional neural network trained in a supervised manner on condensate fractions as labels, which were determined from corresponding rapid ramp images. After successful training, this network architecture allows us to predict the condensate fraction directly from newly acquired time-of-flight images without the need for magnetic field ramps, such as the rapid ramp. The deep convolutional neural network captures subtle features of pairing in the momentum distribution, which are typically too faint for conventional fitting methods, demonstrating the enhanced image recognition capabilities provided by machine learning.

The temperature and density at the predicted onset of superfluidity is then gauged by performing an inverse Abel transformation on corresponding in-situ images, from which the temperature is inferred via fits to the virial expansion of the equation of state. From the resulting critical temperature, we are able to map a phase diagram in the strongly interacting regime of the BEC-BCS crossover. Our results show a striking resemblance to the extended GMB theory in [19], and, for the first time, confirm an increase of the critical temperature from the BCS limit, beyond the unitarity point, with a maximum on the BEC side.

Despite the network's ability to infer condensation directly from time-of-flight images, its inner workings are often obscured by the complexity of the numerous entangled non-linear neuron activations. However, by using the backpropagation-based *DeepLift* library [200], we

have found evidence that the neural network learns a physics-based understanding of pairing and feature extraction for predicting the condensate fraction. This establishes the deep convolutional neural network as a powerful tool for extracting features from atom images, where conventional fitting methods are impractical.

Despite its capabilities, the deep convolutional neural network still requires corresponding condensate fraction labels obtained through the rapid ramp. This becomes particularly problematic when detecting the onset of superfluidity, where the condensate fraction approaches zero. Since these labels are derived from bimodal fits to rapid ramp images, they are susceptible to a low signal-to-noise ratio, possibly affecting predictions close to the onset.

To address this issue, we have introduced a second approach using a deep convolutional autoencoder neural network, an unsupervised machine learning technique. Unlike the first method, the autoencoder does not require labels and is trained (in a supervised manner) on the time-of-flight images with the goal of encoding and decoding of the images to approximate a unity function. Once successfully trained, the encoder efficiently maps input images into a low-dimensional latent space, compressing the significant features of the input data into coordinates of this latent space.

In the latent space, we have found that the data organises itself by temperature and interaction, even though the network is provided no information beyond the time-of-flight images. By analysing the structure of the data in the latent space, we have extracted structural features near the expected superfluid critical temperature, as determined by the first method. Using the same thermometry as before, we have mapped a second phase diagram of the superfluid critical temperature in the BEC-BCS crossover, independently determined without the need for rapid ramp measurements.

As this independent measurement aligns with the phase diagram produced by the supervised neural network, it confirms the validity of our methodology and further supports the observed maximum in the critical temperature between the unitarity point and the BEC limit.

Studying the Higgs mode in the BEC-BCS crossover

In Chapter 6, we have explored the response of the Higgs mode in the BEC-BCS crossover. To this end, we employed two distinct excitation methods on a superfluid Fermi gas within the BEC-BCS crossover, monitoring the gas's response after excitation. Due to the inhomogeneity of our trapping potential, the expected response is highly non-trivial. We attempt to model the dynamics by developing a local density approximation of local Higgs oscillators within the inhomogeneous trap, which contribute to the total spectrum of the inhomogeneous superfluid sample. The model incorporates the measured density distribution, temperature, natural linewidth of the Higgs mode, and the distribution of Cooper pairs in the inhomogeneous condensate fraction, providing us with an extended theoretical framework for comparison with our experimental results.

The first excitation method involves a sudden quench of the interaction parameter by rapidly changing the magnetic field faster than the system's fastest timescale, the Fermi time \hbar/E_F . Immediately after the quench, the order parameter is out of equilibrium with the new interaction parameter, and starts to oscillate. We observe these oscillations in the order parameter as oscillations in the condensate fraction, which we use as an accessible proxy. By fitting the data with a deterministic, phenomenological model, we extract the

oscillation frequency and damping rate. As expected, we observe discernible but damped oscillations for quenches on the BCS side of the crossover. However, as we approach the BEC side, the oscillations become less discernible and eventually undetectable. We account this to the vanishing particle-hole symmetry approaching the BEC side, rendering the Higgs mode unstable.

In comparison with the Fourier-transformed local density approximation model, the observed oscillation frequencies are much slower – by more than a factor of two. Despite this, the damping rates observed in the experiment and predicted by the local density approximation model show strong agreement, suggesting that the significant decay stems from dephasing in the inhomogeneous superfluid. While we have discussed potential reasons for the large discrepancy in oscillation frequency, the developed model could not fully account for this deviation.

In the second approach, we have probed the Higgs mode using a spectroscopic measurement via parametric excitation. To this end, we have tracked the decay rate of the condensate fraction after modulating the magnetic field with a frequency ν_{mod} and amplitude A_{mod} . For frequencies near the expected centre frequency of the local density approximation model, we have observed a peak in the decay rate. We have determined the centre position and width of this peak across the BEC-BCS crossover. Unlike the first method, the observed centre frequency matches well with the model’s predicted centre frequency. However, the observed peak width is significantly narrower than predicted by the model, instead agreeing with a recent measurement of the Higgs mode in an effectively homogenous system [75]. This suggests that while the inhomogeneity of the trap lowers the observed centre Higgs frequency, it does not lead to a significant broadening of the Higgs mode.

Additionally, the spectroscopic approach has also allowed us to investigate the temperature dependence of the Higgs mode. Contrary to expectation, we observe no change in the oscillation frequency with temperature, which, aligns with a recent measurement of the Higgs mode in an effectively homogeneous system [75]. This unexpected result warrants further investigation.

Cooling by dimer dissociation

In Chapter 7, we have developed and demonstrated a novel cooling technique based on the selective dissociation and removal of high-energy dimers from an inhomogeneous trapping potential. This cooling method operates on the BEC side of the BEC-BCS crossover, where the Feshbach resonance supports a two body-bound state, and dimers can be dissociated by modulating the magnetic field with frequencies higher than the corresponding bound state energy. Dissociating a dimer in a pure dimer cloud is connected to an increase in the repulsive mean-field energy, which results in the dissociation threshold energy to experience a density-proportional shift from the bound state energy. This shift allows for a selective dissociation and subsequent removal of predominantly high-energy dimers, effectively cooling the remaining sample via evaporative cooling.

We have characterised the cooling technique by measuring the phase-space density, either directly or through the number of condensed atoms as a proxy, probed after an additional thermalisation time subsequent to the modulation of the magnetic field. For modulation frequencies greater than the bound state energy, we have observed an increase in phase-space

density. The corresponding temperature reduction has been explained using the step model of evaporation, applied to a simulated dimer sample with the measured temperature, atom number and trap frequencies.

For a more quantitative description of the cooling technique, we have determined the cooling efficiency for samples across the BEC side of the BEC-BCS crossover. In general, the cooling technique is most effective and efficient for samples with the highest bound state energy, furthest on the BEC side. We have measured efficiencies of up to $\gamma = 4$ in a gravity-tilted trap and up to $\gamma = 6$ in a gravity-compensated trap, matching and even exceeding the highest reported efficiencies in forced evaporation implementations near a Feshbach resonance [85, 87, 91–99]. As we approach unitarity, the bound state energy decreases, causing the cooling technique to become less efficient, eventually approaching $\gamma = 1$. For bound state energies below $E_B \sim 30$ kHz no cooling has been detected. We have attributed the decreasing efficiency towards unitarity to two effects aligning with our proposed cooling mechanism: a gradual decrease in bound state energy E_B and a shift from two-body dynamics to strong many-body interactions.

Finally, we have investigated the timescale of the cooling technique, which scales as $t_{\text{mod}} \propto 1/\sqrt{A_{\text{mod}}}$, with A_{mod} being the modulation amplitude. For our apparatus, the fastest cooling step is limited by the maximum achievable modulation amplitude, resulting in a modulation time of $t_{\text{mod}} = 9.3$ ms and a subsequent $1/e$ thermalisation time of 20 ms.

It is important to note that the selective removal operates through magnetic field modulation and, unlike radio-frequency-transitions (RF), does not couple to a third spin state. This prevents strong three-body losses during evaporation. Moreover, the cooling process does not require alterations to the trapping potential, simplifying the process of maintaining the efficient runaway regime. In conclusion, this cooling technique should be applicable for a wide range of systems, achieving high efficiencies.

8.2 Outlook

In addition to the results presented in Section 8.1, the experimental apparatus offers opportunities for future endeavours closely related to the studies in this thesis. Here, we outline three potential projects that can be realised with the current setup, requiring minimal modifications to the system.

8.2.1 Probing the Higgs mode in a homogeneous Fermi gas

In Chapter 6, we have studied the response of the Higgs mode in the BEC-BCS crossover for the inhomogeneous sample produced by our apparatus. As discussed in Section 6.1, the inhomogeneity of the sample necessitates extended models to describe the observed oscillations. We have attempted to address this by modelling the Higgs mode spectrum through a local density approximation of local Higgs oscillators. Nonetheless, probing the Higgs mode in a homogeneous system would allow for a more direct investigation, as effectively demonstrated in [75].

To directly compare our results with those of an equivalent homogeneous system, it would be advantageous to repeat our measurements after preparing the Fermi gas in such a system. In our current setup, inhomogeneity arises from the crossed optical dipole trap (see

Section 3.2.6). In Appendix A, we have presented a suitable setup for creating a homogenous box trap potential within our setup.

This setup is already available and can be implemented in future measurements. Of course, analysing results from a homogeneous Fermi gas would introduce its own set of challenges, such as a modified thermometry due to the altered potential.

8.2.2 The Higgs mode in the crossover from 3D to 2D

In Section 2.3.4, we have introduced the Higgs mode as a consequence of spontaneous symmetry breaking upon entering the superfluid state. When transitioning from a 3D to a 2D system, the superfluid phase transition is replaced by the Berezinskii-Kosterlitz-Thouless transition [231], which describes a superfluid state without long-range order or symmetry breaking [232]. Despite ongoing debate, some predictions still suggest that the Higgs mode persists in 2D systems [233, 234].

The Fermi gas in our experimental apparatus offers an ideal platform to probe the Higgs mode within the transition from a 3D to a 2D system. For this purpose, a setup similar to that described in Appendix A can be used to create a blue-detuned optical dipole trap. This trap employs the 532 nm TEM01 mode along one spatial direction with a 589 nm homogenous ring potential in the remaining plane, effectively suppressing dynamics in one spatial direction and rendering the system two-dimensional. By combining this setup with the custom-built small magnetic field coil, we can probe excitations while transitioning the gas from 3D to 2D.

8.2.3 Measuring the pseudogap

In Chapter 5, we have measured the phase diagram of the superfluid critical temperature T_C in the BEC-BCS crossover. For temperatures below the superfluid critical temperature $T < T_C$, a superfluid condensate of paired atoms forms. As discussed in Section 2.3, the critical temperature in the BCS limit coincides with the onset of pairing $T^* = T_C$. However, in the BEC limit, preformed pairs exist at much higher temperatures than the critical temperature $T^* \gg T_C$. Thus, in the strongly interacting regime of the BEC-BCS crossover, thermal pairs can exist at temperatures above the critical temperature T_C . We have discussed such pairs in a pseudogap as a possible explanation of the missing temperature dependence of the Higgs frequency in Fig. 6.12. The existence of preformed pairs could also influence the number of condensed pairs from a rapid ramp measurement (refer to Section 4.3).

The key question is whether the strong interaction between these pairs lead to a pseudogap state at a temperature T between the onset of pairing and the onset of superfluidity $T_C < T < T^*$. The interpretation and existence of such a pseudogap state in a strongly interacting Fermi gas is highly debated [222, 235]. In some theories, the pseudogap phase is interpreted as a regime where the gas exhibits a BCS-like dispersion relation of incoherent pairs, as described in Eq. (2.19), but with the gap parameter Δ no longer describing the order parameter of a macroscopically occupied coherent state [223]. In previous works, the pseudogap has been experimentally investigated by observing back-bending at the Fermi momentum k_F in the dispersion relation [223, 236], or theoretical discussed in terms of its effect on gas properties, such as pairing correlations [222].

In our apparatus, the pseudogap can be probed using an approach similar to the measurements in Section 6.3.1 and Section 7.3. One naive approach would be to modulate the magnetic field at a frequency ν_{mod} high enough that it can only couple to the pseudogap parameter Δ . This is done for samples at various initial temperatures T_i/T_F , set by a time-dependent recompression of the trapping potential for a time of t_{heat} , as described in Section 5.3.1. If the gas is above a temperature that supports a pseudogap state, the modulation should not couple to any relevant modes, and the gas should be less perturbed. We have tested this approach for a strongly interacting Fermi gas at unitarity, as shown in Fig. 8.1, by measuring the temperature before and after the perturbation (refer to Section 4.2).

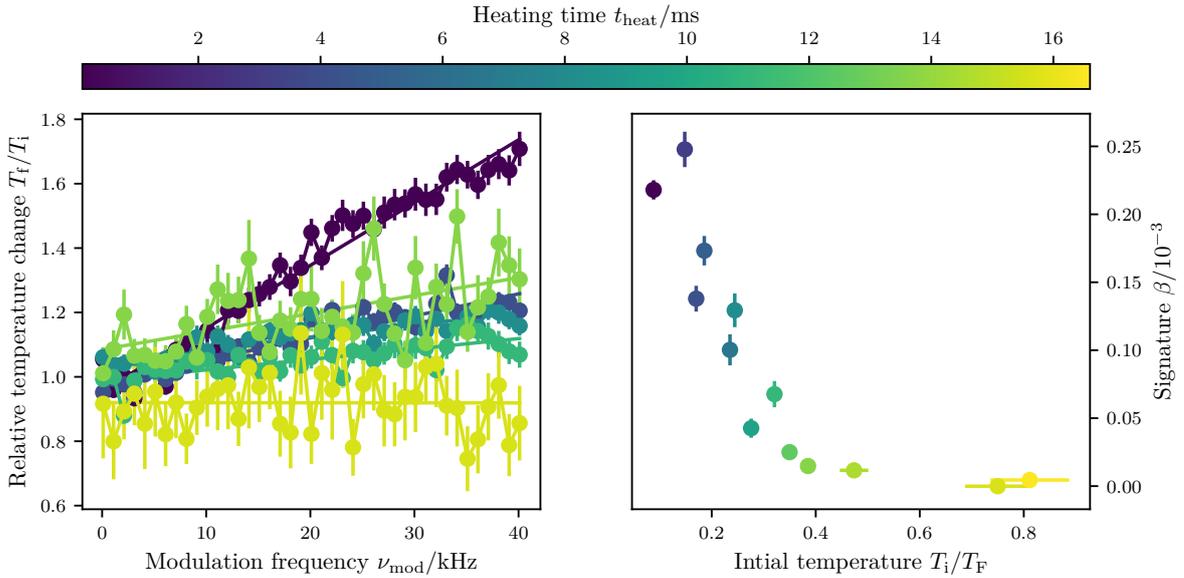


Figure 8.1: Investigating the pseudogap for a Fermi gas at unitarity. We have measured the response to a magnetic field modulation at maximum modulation amplitude and for varying modulation frequencies ν_{mod} . This measurement is performed for samples at various initial temperatures T_i/T_F , which depends on the heating time t_{heat} before the modulation, as described in Section 5.3.1. **Left:** For the coldest samples with the lowest heating time t_{heat} , the sample temperature T_f/T_i changes most drastically with an increasing modulation frequency ν_{mod} . **Right:** For each heating time t_{heat} , we show the signature $\beta = m/\chi^2$, extracted from a linear fit to the data in the **left** panel. Here, m describes the slope and χ^2 represents the fit quality.

For the coldest initial samples at $T_i/T_F < T_C/T_F = 0.167(13)$ [165] a superfluid still exists, and strong coupling results in significant heating. As we increase the initial temperature T_i/T_F , the heating effect diminishes. Since the temperature determination becomes noisier due to the stronger atom loss with higher initial temperature, we define a signature $\beta = m/\chi^2$ for keeping track of the heating rate. The parameter m is the slope extracted from a linear fit to the data, while χ^2 reflects the quality of the fit. This method prevents the slope from being overestimated due to noise in the temperature measurements.

The decreasing β with increasing initial temperature T_i/T_F suggests a diminishing coupling strength to the sample. Our data indicate a change in trend between temperatures of $T_i/T_F = 0.3$ and 0.4 , which agrees with earlier studies showing a possible pseudogap transition

around $T_i/T_F < 0.35$ [223] and $T_i/T_F = 0.29$ [236].

However, it is important to note that the diminishing signature could also be explained by the decreasing density with increasing initial temperature, which slows down thermalisation and heating. Further research is necessary to better understand the observed dynamics and to clarify a connection to the possible existence of a pseudogap state. A good starting point would be to track the heating behaviour across the strongly interacting regime of the BEC-BCS crossover and to eliminate crosstalk of the diminishing density.

Box trap setup

In our apparatus, a homogeneous system can be realised using a blue-detuned optical box potential. To this end, my colleague Andreas Kell has already implemented an axicon ring beam setup for 589 nm [237], which defines the potential walls in the xy -plane along the vertical z -direction. Originally, the box trap was designed to be closed in z -direction by end-caps formed by a 532 nm TEM01 mode [237]. The targeted dimensions for the cylindrical box are approximately $\sim 80 \mu\text{m}$ for the ring radius and the vertical end-cap distance. A larger end-cap distance would introduce stronger inhomogeneities due to the remaining curvature of the magnetic field. Although this setup would suffice in trapping atoms, the TEM01 mode would introduce a harmonic component along the z -direction, leading to further inhomogeneities.

Light sheet setup

During the works on this thesis, we have integrated an alternative solution for the potential end-caps in the form of two light sheet beams, referred to as light sheets henceforth. The implemented setup is sketched in Fig. A.1, and makes use of the same 532 nm light as well.

The setup is constrained by the final lens, which has a focal length of $f = 160 \text{ mm}$ and focuses the beam onto the atoms. Since the light is shared with the plug beam setup (refer to Section 3.2.5), the initial q beam parameters are fixed to $q_V = (-3.39 + 0.723i) \text{ m}$ and $q_H = (-5.55 + 2.240i) \text{ m}$ for the vertical and horizontal direction, respectively [237].

The light sheets are created using a Wollaston prism, which splits the beam into two beams with a 1° separation. The light sheets require the beams to focus vertically at the atom position, while the horizontal direction remains relatively collimated and wide. Several cylindrical lenses are placed around the setup to achieve this. We have simulated the proposed setup using ray transfer matrix analysis [160] and optimised the lens distances to achieve the desired light sheet dimensions. The simulated sheet distance as well and vertical w_V and horizontal w_H waists are shown in Fig. A.2.

At the atom position, the simulation yields a light sheet distance of $79.6 \mu\text{m}$, with a vertical waist of $w_V = 7.2 \mu\text{m}$ and a horizontal waist of $w_H = 953.2 \mu\text{m}$. The corresponding Rayleigh ranges are $z_{R,V} = 0.3 \text{ mm}$ for the vertical direction and $z_{R,H} = 5.4 \text{ m}$ for the horizontal

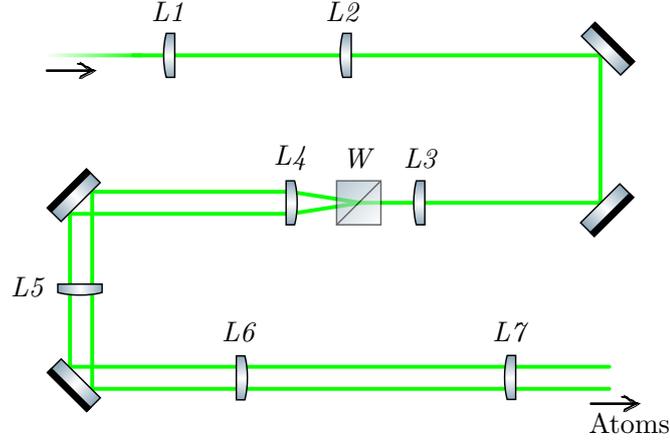


Figure A.1: Setup for the creation of light sheets. The setup for creating light sheets is shown, with the beam direction indicated by arrows. The optics include seven plano-convex lenses $L1-L7$, one Wollaston prism W , and four mirrors (not labelled). All optics are $\varnothing 1$ in. The focal lengths and orientations of the lenses are: $f_1 = 1000$ mm (cyl. vert.), $f_2 = 100$ mm (cyl. hori.), $f_3 = 50$ mm (cyl. vert.), $f_4 = 13.7$ mm (cyl. vert.), $f_5 = 200$ mm (cyl. hori.), $f_6 = 500$ mm (sph.) and $f_7 = 160$ mm (sph.). The Wollaston prism (Thorlabs WPQ10) has a beam separation angle of 1° . Sketch uses components from [238].

direction. This geometry is suitable for the desired box trap end-caps and has already been implemented into the apparatus. The light sheets have been roughly aligned to the atom position, and a picture of the light sheets taken with ANDOR2 (refer to Section 4.1.1) along the atom imaging path confirms the desired geometry.

The next steps involve characterising the box trap and loading it with atoms. Afterwards, the transfer from the crossed optical dipole trap to the box trap has to be optimised to achieve good phase-space densities. Due to the larger optical access required for the ring beam from the top of the chamber, the currently installed small magnetic field coil (refer to Section 4.5) must be replaced with a similar coil of larger inner diameter. A suitable coil has already been designed and must be integrated before adding the ring beam.

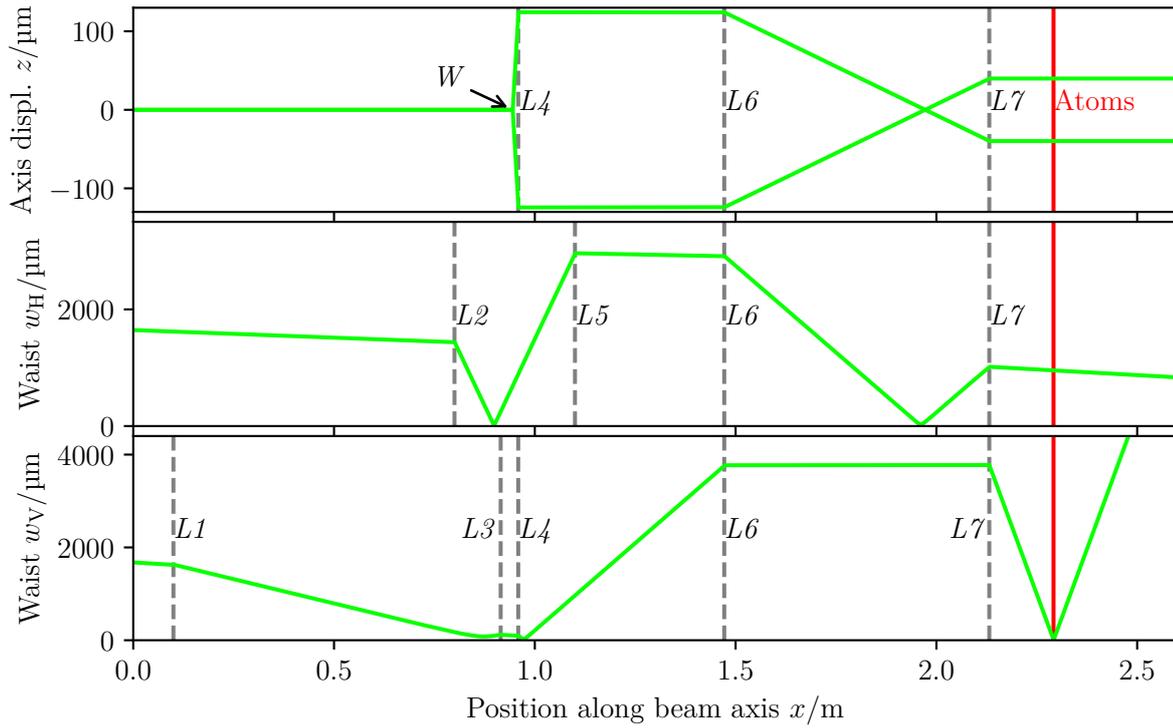


Figure A.2: Ray transfer matrix analysis of the light sheet setup. The setup from Fig. A.1 is simulated via ray transfer matrix analysis to optimise component positions (dashed vertical lines) such that the light sheets achieve the desired geometry at the atom position (red vertical line). **Top:** Axis displacement of the two beams forming the light sheets. The single beam is split into two beams at the Wollaston prism (W). **Centre and bottom:** Waist in the horizontal w_H and vertical w_V directions, respectively.

Special functions

Special functions used in the thesis are specified here.

Skewed Gaussian function

The skewed Gaussian function is defined as

$$f(x) = 2A \cdot \phi(x) \cdot \Phi(x), \quad (\text{B.1})$$

with

$$\phi(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(x-x_0)^2}{2\sigma^2}\right], \quad (\text{B.2})$$

and

$$\Phi(x) = 1/2 \left[1 + \operatorname{erf}\left(\alpha \cdot \frac{x-x_0}{\sqrt{2}\sigma}\right) \right]. \quad (\text{B.3})$$

Here, A is the amplitude, x_0 is the centre position, σ is the width, α is the skeweness and $\operatorname{erf}(x) := \int_{-\infty}^x \phi(t)dt$ denotes the error function. This function is used in Section 7.3.3.

Bibliography

- [1] H. K. Onnes, *The Nobel Prize in Physics 1913*, 1913, URL: <https://www.nobelprize.org/prizes/physics/1913/onnes/lecture/>.
- [2] J. Hirsch, M. Maple and F. Marsiglio, *Superconducting materials classes: Introduction and overview*, *Physica C: Superconductivity and its Applications* **514** (2015) 1, URL: <https://www.sciencedirect.com/science/article/pii/S0921453415000933>.
- [3] J. Liu et al., *World record 32.35 tesla direct-current magnetic field generated with an all-superconducting magnet*, *Superconductor Science and Technology* **33** (2020) 03LT01, URL: <https://dx.doi.org/10.1088/1361-6668/ab714e>.
- [4] M. Parizh, Y. Lvovsky and M. Sumption, *Conductors for commercial MRI magnets beyond NbTi: requirements and challenges*, *Superconductor science & technology* **30** (2016) 014007, URL: <https://pmc.ncbi.nlm.nih.gov/articles/PMC5472374/>.
- [5] Y. Y. Gao, M. A. Rol, S. Touzard and C. Wang, *Practical Guide for Building Superconducting Quantum Devices*, *PRX Quantum* **2** (2021) 040202, URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.040202>.
- [6] P. Krantz et al., *A quantum engineer's guide to superconducting qubits*, *Applied Physics Reviews* **6** (2019) 021318, URL: <https://doi.org/10.1063/1.5089550>.
- [7] M. Y. Kagan, *Modern trends in Superconductivity and Superfluidity*, vol. 874, *Lecture Notes in Physics*, Springer, 2013, URL: <http://link.springer.com/10.1007/978-94-007-6961-8>.
- [8] W. Ketterle and M. W. Zwierlein, *Making, probing and understanding ultracold Fermi gases*, *La Rivista del Nuovo Cimento* **31** (2008) 247, URL: <http://arxiv.org/abs/0801.2500>.
- [9] J. F. Allen and A. D. Misener, *Flow Phenomena in Liquid Helium II*, *Nature* **142** (1938) 643, URL: <https://www.nature.com/articles/142643a0>.

- [10] S. Balibar, *The Discovery of Superfluidity*, Journal of Low Temperature Physics **146** (2007) 441, URL: <https://doi.org/10.1007/s10909-006-9276-7>.
- [11] A. Einstein, *Quantentheorie des einatomigen idealen Gases. Zweite Abhandlung*, John Wiley & Sons, Ltd, 2005, URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/3527608958.ch28>.
- [12] L. D. Landau, *The Nobel Prize in Physics 1962*, 1962, URL: <https://www.nobelprize.org/prizes/physics/1962/summary/>.
- [13] J. Bardeen, L. N. Cooper and J. R. Schrieffer, *Theory of Superconductivity*, Physical Review **108** (1957) 1175, URL: <https://link.aps.org/doi/10.1103/PhysRev.108.1175>.
- [14] L. N. Cooper, *Bound Electron Pairs in a Degenerate Fermi Gas*, Physical Review **104** (1956) 1189, URL: <https://link.aps.org/doi/10.1103/PhysRev.104.1189>.
- [15] A. J. Leggett, *Diatomic molecules and cooper pairs*, Modern Trends in the Theory of Condensed Matter (1980) 13, URL: <https://link.springer.com/chapter/10.1007/BFb0120125>.
- [16] P. Nozières and S. Schmitt-Rink, *Bose condensation in an attractive fermion gas: From weak to strong coupling superconductivity*, Journal of Low Temperature Physics **59** (1985) 195, URL: <https://doi.org/10.1007/BF00683774>.
- [17] C. H. Schunck, Y.-i. Shin, A. Schirotzek and W. Ketterle, *Determination of the fermion pair size in a resonantly interacting superfluid*, Nature **454** (2008) 739, URL: <https://www.nature.com/articles/nature07176>.
- [18] W. Zwerger, ed., *The BCS-BEC Crossover and the Unitary Fermi Gas*, vol. 836, Lecture Notes in Physics, Springer, 2012, URL: <https://link.springer.com/10.1007/978-3-642-21978-8>.
- [19] L. Pisani, A. Perali, P. Pieri and G. C. Strinati, *Entanglement between pairing and screening in the Gorkov-Melik-Barkhudarov correction to the critical temperature throughout the BCS-BEC crossover*, Physical Review B **97** (2018) 014528, URL: <https://link.aps.org/doi/10.1103/PhysRevB.97.014528>.
- [20] S. Pilati, S. Giorgini and N. Prokof'ev, *Critical Temperature of Interacting Bose Gases in Two and Three Dimensions*, Physical Review Letters **100** (2008) 140405, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.100.140405>.
- [21] R. P. Feynman, *Simulating physics with computers*, International Journal of Theoretical Physics **21** (1982) 467, URL: <https://doi.org/10.1007/BF02650179>.
- [22] S. Lloyd, *Universal Quantum Simulators*, Science **273** (1996) 1073, URL: <https://www.science.org/doi/10.1126/science.273.5278.1073>.

-
- [23] H. Lin et al., *Real-space BCS-BEC crossover in FeSe monolayers*, Physical Review B **107** (2023) 104517, URL: <https://link.aps.org/doi/10.1103/PhysRevB.107.104517>.
- [24] T. W. Hänsch and A. L. Schawlow, *Cooling of gases by laser radiation*, Optics Communications **13** (1975) 68, URL: <https://www.sciencedirect.com/science/article/pii/0030401875901595>.
- [25] S. Chu, *Nobel Lecture: The manipulation of neutral particles*, Reviews of Modern Physics **70** (1998) 685, URL: <https://link.aps.org/doi/10.1103/RevModPhys.70.685>.
- [26] C. N. Cohen-Tannoudji, *Nobel Lecture: Manipulating atoms with photons*, Reviews of Modern Physics **70** (1998) 707, URL: <https://link.aps.org/doi/10.1103/RevModPhys.70.707>.
- [27] W. D. Phillips, *Nobel Lecture: Laser cooling and trapping of neutral atoms*, Reviews of Modern Physics **70** (1998) 721, URL: <https://link.aps.org/doi/10.1103/RevModPhys.70.721>.
- [28] H. Feshbach, *Unified theory of nuclear reactions*, Annals of Physics **5** (1958) 357, URL: <https://www.sciencedirect.com/science/article/pii/0003491658900071>.
- [29] S. Inouye et al., *Observation of Feshbach resonances in a Bose-Einstein condensate*, Nature **392** (1998) 151, URL: <https://www.nature.com/articles/32354>.
- [30] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman and E. A. Cornell, *Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor*, Science **269** (1995) 198, URL: <https://www.science.org/doi/10.1126/science.269.5221.198>.
- [31] B. DeMarco and D. S. Jin, *Onset of Fermi Degeneracy in a Trapped Atomic Gas*, Science **285** (1999) 1703, URL: <https://www.science.org/doi/10.1126/science.285.5434.1703>.
- [32] J. Cubizolles, T. Bourdel, S. J. J. M. F. Kokkelmans, G. V. Shlyapnikov and C. Salomon, *Production of Long-Lived Ultracold Li₂ Molecules from a Fermi Gas*, Physical Review Letters **91** (2003) 240401, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.91.240401>.
- [33] C. A. Regal, C. Ticknor, J. L. Bohn and D. S. Jin, *Creation of ultracold molecules from a Fermi gas of atoms*, Nature **424** (2003) 47, URL: <https://www.nature.com/articles/nature01738>.
- [34] S. Jochim et al., *Bose-Einstein Condensation of Molecules*, Science **302** (2003) 2101, URL: <https://www.science.org/doi/full/10.1126/science.1093280>.
- [35] M. Greiner, C. A. Regal and D. S. Jin, *Emergence of a molecular Bose-Einstein condensate from a Fermi gas*, Nature **426** (2003) 537, URL: <http://www.nature.com/articles/nature02199>.
- [36] M. W. Zwierlein et al., *Observation of Bose-Einstein Condensation of Molecules*, Physical Review Letters **91** (2003) 250401, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.91.250401>.

- [37] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch and I. Bloch, *Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms*, Nature **415** (2002) 39, URL: <https://www.nature.com/articles/415039a>.
- [38] C. A. Regal, M. Greiner and D. S. Jin, *Observation of Resonance Condensation of Fermionic Atom Pairs*, Physical Review Letters **92** (2004) 040403, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.92.040403>.
- [39] J. Kinast, S. L. Hemmer, M. E. Gehm, A. Turlapov and J. E. Thomas, *Evidence for Superfluidity in a Resonantly Interacting Fermi Gas*, Physical Review Letters **92** (2004) 150402, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.92.150402>.
- [40] A. Behrle, *Driving a Strongly Interacting Superfluid out of Equilibrium*, PhD thesis: University of Bonn, 2017.
- [41] T. J. Harrison, *Measuring the Gap and Investigating Non-equilibrium in the BEC-BCS Crossover*, PhD thesis: University of Bonn, 2017.
- [42] P. Arnold and G. Moore, *BEC Transition Temperature of a Dilute Homogeneous Imperfect Bose Gas*, Physical Review Letters **87** (2001) 120401, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.87.120401>.
- [43] V. A. Kashurnikov, N. V. Prokof'ev and B. V. Svistunov, *Critical Temperature Shift in Weakly Interacting Bose Gas*, Physical Review Letters **87** (2001) 120402, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.87.120402>.
- [44] L. P. Gor'Kov and T. K. Melik-Barkhudarov, *Contribution to the Theory of Superfluidity in an Imperfect Fermi Gas*, JETP **13** (1961) 1018, URL: <http://www.jetp.ras.ru/cgi-bin/e/index/e/13/5/p1018?a=list>.
- [45] R. Haussmann, W. Rantner, S. Cerrito and W. Zwerger, *Thermodynamics of the BCS-BEC crossover*, Physical Review A **75** (2007) 023610, URL: <https://link.aps.org/doi/10.1103/PhysRevA.75.023610>.
- [46] E. Burovski, N. Prokof'ev, B. Svistunov and M. Troyer, *Critical Temperature and Thermodynamics of Attractive Fermions at Unitarity*, Physical Review Letters **96** (2006) 160402, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.96.160402>.
- [47] A. Bulgac, J. E. Drut and P. Magierski, *Quantum Monte Carlo simulations of the BCS-BEC crossover at finite temperature*, Physical Review A **78** (2008) 023625, URL: <https://link.aps.org/doi/10.1103/PhysRevA.78.023625>.

-
- [48] W. Ketterle, D. S. Durfee and D. M. Stamper-Kurn, *Making, probing and understanding Bose-Einstein condensates*, 1999, arXiv: cond-mat/9904034 [cond-mat], URL: <http://arxiv.org/abs/cond-mat/9904034>.
- [49] M. W. Zwierlein et al., *Condensation of Pairs of Fermionic Atoms near a Feshbach Resonance*, Physical Review Letters **92** (2004) 120403, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.92.120403>.
- [50] A. Perali, P. Pieri and G. C. Strinati, *Extracting the Condensate Density from Projection Experiments with Fermi Gases*, Physical Review Letters **95** (2005) 010407, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.95.010407>.
- [51] E. Altman and A. Vishwanath, *Dynamic Projection on Feshbach Molecules: A Probe of Pairing and Phase Fluctuations*, Physical Review Letters **95** (2005) 110404, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.95.110404>.
- [52] E. A. Yuzbashyan, B. L. Altshuler, V. B. Kuznetsov and V. Z. Enolskii, *Nonequilibrium cooper pairing in the nonadiabatic regime*, Physical Review B **72** (2005) 220503, URL: <https://link.aps.org/doi/10.1103/PhysRevB.72.220503>.
- [53] Y. Liu et al., *Summary of ChatGPT-Related Research and Perspective Towards the Future of Large Language Models*, 2023, arXiv: 2304.01852 [cs.LG], URL: <http://arxiv.org/abs/2304.01852>.
- [54] M. A. Kramer, *Nonlinear principal component analysis using autoassociative neural networks*, AIChE Journal **37** (1991) 233, URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/aic.690370209>.
- [55] A. Blattmann, R. Rombach, K. Oktay, J. Müller and B. Ommer, *Semi-Parametric Neural Image Synthesis*, 2022, arXiv: 2204.11824 [cs.CV], URL: <http://arxiv.org/abs/2204.11824>.
- [56] A. Dawid et al., *Modern applications of machine learning in quantum sciences*, 2023, arXiv: 2204.04198 [quant-ph], URL: <http://arxiv.org/abs/2204.04198>.
- [57] M. Krenn, J. Landgraf, T. Foesel and F. Marquardt, *Artificial intelligence and machine learning for quantum technologies*, Physical Review A **107** (2023) 010101, URL: <https://link.aps.org/doi/10.1103/PhysRevA.107.010101>.
- [58] J. J. Hopfield and G. E. Hinton, *The Nobel Prize in Physics 2024*, 2024, URL: <https://www.nobelprize.org/prizes/physics/2024/summary/>.
- [59] N. Milson et al., *High-dimensional reinforcement learning for optimization and control of ultracold quantum gases*, Machine Learning: Science and Technology **4** (2023) 045057, URL: <https://dx.doi.org/10.1088/2632-2153/ad1437>.

- [60] P. Kobel, M. Link and M. Köhl, *Exponentially improved detection and correction of errors in experimental systems using neural networks*, 2020, arXiv: 2005.09119 [quant-gas], URL: <http://arxiv.org/abs/2005.09119>.
- [61] E. P. L. van Nieuwenburg, Y.-H. Liu and S. D. Huber, *Learning phase transitions by confusion*, Nature Physics **13** (2017) 435, URL: <https://www.nature.com/articles/nphys4037>.
- [62] L. Wang, *Discovering phase transitions with unsupervised learning*, Physical Review B **94** (2016) 195105, URL: <https://link.aps.org/doi/10.1103/PhysRevB.94.195105>.
- [63] C. Alexandrou, A. Athenodorou, C. Chrysostomou and S. Paul, *The critical temperature of the 2D-Ising model through deep learning autoencoders*, The European Physical Journal B **93** (2020) 226, URL: <http://link.springer.com/10.1140/epjb/e2020-100506-5>.
- [64] W. Hu, R. R. P. Singh and R. T. Scalettar, *Discovering phases, phase transitions, and crossovers through unsupervised machine learning: A critical examination*, Physical Review E **95** (2017) 062122, URL: <https://link.aps.org/doi/10.1103/PhysRevE.95.062122>.
- [65] B. S. Rem et al., *Identifying Quantum Phase Transitions using Artificial Neural Networks on Experimental Data*, Nature Physics **15** (2019) 917, URL: <https://www.nature.com/articles/s41567-019-0554-0>.
- [66] N. Käming et al., *Unsupervised machine learning of topological phase transitions from experimental data*, Machine Learning: Science and Technology **2** (2021) 035037, URL: <https://iopscience.iop.org/article/10.1088/2632-2153/abffe7>.
- [67] M. Link et al., *Machine Learning the Phase Diagram of a Strongly Interacting Fermi Gas*, Physical Review Letters **130** (2023) 203401, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.203401>.
- [68] D. Eberz et al., *Detecting the phase transition in a strongly interacting Fermi gas by unsupervised machine learning*, Physical Review A **108** (2023) 063303, URL: <https://link.aps.org/doi/10.1103/PhysRevA.108.063303>.
- [69] L. Salasnich, N. Manini and A. Parola, *Condensate fraction of a Fermi gas in the BCS-BEC crossover*, Physical Review A **72** (2005) 023621, URL: <https://link.aps.org/doi/10.1103/PhysRevA.72.023621>.
- [70] D. Pekker and C. M. Varma, *Amplitude/Higgs Modes in Condensed Matter Physics*, Annual Review of Condensed Matter Physics **6** (2015) 269, URL: <https://www.annualreviews.org/content/journals/10.1146/annurev-conmatphys-031214-014350>.
- [71] Y. Nambu, *Quasi-Particles and Gauge Invariance in the Theory of Superconductivity*, Physical Review **117** (1960) 648, URL: <https://link.aps.org/doi/10.1103/PhysRev.117.648>.

-
- [72] J. Goldstone, A. Salam and S. Weinberg, *Broken Symmetries*, Physical Review **127** (1962) 965,
URL: <https://link.aps.org/doi/10.1103/PhysRev.127.965>.
- [73] A. F. Volkov and S. M. Kogan,
Collisionless relaxation of the energy gap in superconductors, JETP **38** (1974) 1018,
URL: <http://jetp.ras.ru/cgi-bin/e/index/e/38/5/p1018?a=list>.
- [74] R. Matsunaga et al., *Higgs Amplitude Mode in the BCS Superconductors Nb_{1-x}Ti_xN Induced by Terahertz Pulse Excitation*, Physical Review Letters **111** (2013) 057002,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.111.057002>.
- [75] P. Dyke et al., *Higgs Oscillations in a Unitary Fermi Superfluid*, Physical Review Letters **132** (2024) 223402,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.132.223402>.
- [76] S. Hannibal et al., *Quench dynamics of an ultracold Fermi gas in the BCS regime: Spectral properties and confinement-induced breakdown of the Higgs mode*, Physical Review A **91** (2015) 043630,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.91.043630>.
- [77] S. Hannibal, P. Kettmann, M. D. Croitoru, V. M. Axt and T. Kuhn,
Persistent oscillations of the order parameter and interaction quench phase diagram for a confined Bardeen-Cooper-Schrieffer Fermi gas, Physical Review A **98** (2018) 053605,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.98.053605>.
- [78] A. Behrle et al., *Higgs mode in a strongly interacting fermionic superfluid*, Nature Physics **14** (2018) 781,
URL: <https://www.nature.com/articles/s41567-018-0128-6>.
- [79] R. G. Scott, F. Dalfovo, L. P. Pitaevskii and S. Stringari,
Rapid ramps across the BEC-BCS crossover: A route to measuring the superfluid gap, Physical Review A **86** (2012) 053604,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.86.053604>.
- [80] R. Matsunaga et al., *Light-induced collective pseudospin precession resonating with Higgs mode in a superconductor*, Science **345** (2014) 1145,
URL: <https://www.science.org/doi/10.1126/science.1254697>.
- [81] H. P. Ojeda Collado, G. Usaj, C. A. Balseiro, D. H. Zanette and J. Lorenzana,
Emergent parametric resonances and time-crystal phases in driven Bardeen-Cooper-Schrieffer systems, Physical Review Research **3** (2021) L042023,
URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.3.L042023>.
- [82] G. Lyu, K.-T. Xi, S. Yoon, Q. Chen and G. Watanabe,
Exciting the long-lived Higgs mode in superfluid Fermi gases with particle removal, Physical Review A **107** (2023) 023321,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.107.023321>.

- [83] A. Barresi, A. Boulet, G. Wlazłowski and P. Magierski, *Generation and decay of Higgs mode in a strongly interacting Fermi gas*, Scientific Reports **13** (2023) 11285, URL: <https://www.nature.com/articles/s41598-023-38176-9>.
- [84] C. M. Varma, *Higgs Boson in Superconductors*, Journal of Low Temperature Physics **126** (2002) 901, URL: <https://doi.org/10.1023/A:1013890507658>.
- [85] W. Ketterle and N. V. Druten, *Evaporative Cooling of Trapped Atoms*, Academic Press: vol. 37, 1996, URL: <https://linkinghub.elsevier.com/retrieve/pii/S1049250X08601019>.
- [86] K. B. Davis, M.-O. Mewes, M. A. Joffe, M. R. Andrews and W. Ketterle, *Evaporative Cooling of Sodium Atoms*, Physical Review Letters **74** (1995) 5202, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.74.5202>.
- [87] S. R. Granade, M. E. Gehm, K. M. O'Hara and J. E. Thomas, *All-Optical Production of a Degenerate Fermi Gas*, Physical Review Letters **88** (2002) 120405, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.88.120405>.
- [88] M. Sigrist, *Unkonventionelle Supraleitung*, 2006.
- [89] A. Y. Kitaev, *Unpaired Majorana fermions in quantum wires*, Physics-Uspekhi **44** (2001) 131, URL: <https://dx.doi.org/10.1070/1063-7869/44/10S/S29>.
- [90] J. Knolle and R. Moessner, *A Field Guide to Spin Liquids*, Annual Review of Condensed Matter Physics **10** (2019) 451, URL: <https://www.annualreviews.org/content/journals/10.1146/annurev-conmatphys-031218-013401>.
- [91] K. M. O'Hara, *Optical trapping and evaporative cooling of fermionic atoms*, PhD thesis: Duke University, 2000.
- [92] S. Jochim et al., *Pure Gas of Optically Trapped Molecules Created from Fermionic Atoms*, Physical Review Letters **91** (2003) 240402, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.91.240402>.
- [93] T. Bourdel et al., *Experimental Study of the BEC-BCS Crossover Region in Lithium 6*, Physical Review Letters **93** (2004) 050401, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.93.050401>.
- [94] T. Kinoshita, T. Wenger and D. S. Weiss, *All-optical Bose-Einstein condensation using a compressible crossed dipole trap*, Physical Review A **71** (2005) 011602, URL: <https://link.aps.org/doi/10.1103/PhysRevA.71.011602>.

-
- [95] J.-F. Clément et al., *All-optical runaway evaporation to Bose-Einstein condensation*, Physical Review A **79** (2009) 061406,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.79.061406>.
- [96] A. J. Olson, R. J. Niffenegger and Y. P. Chen,
Optimizing the efficiency of evaporative cooling in optical dipole traps, Physical Review A **87** (2013) 053613,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.87.053613>.
- [97] R. Roy, A. Green, R. Bowler and S. Gupta,
Rapid cooling to quantum degeneracy in dynamically shaped atom traps, Physical Review A **93** (2016) 043403,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.93.043403>.
- [98] K. Yamashita, K. Hanasaki, A. Ando, M. Takahama and T. Kinoshita,
All-optical production of a large Bose-Einstein condensate in a double compressible crossed dipole trap, Physical Review A **95** (2017) 013609,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.95.013609>.
- [99] A. Herbst, H. Albers, K. Stolzenberg, S. Bode and D. Schlippert, *Rapid generation of all-optical ^{39}K Bose-Einstein condensates using a low-field Feshbach resonance*, Physical Review A **106** (2022) 043320,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.106.043320>.
- [100] A. Kell, M. Breyer, D. Eberz and M. Köhl, *Exciting the Higgs Mode in a Strongly Interacting Fermi Gas by Interaction Modulation*, Physical Review Letters **133** (2024) 150403,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.133.150403>.
- [101] M. Breyer, D. Eberz, A. Kell and M. Köhl,
Quenching a Fermi superfluid across the BEC-BCS crossover, SciPost Physics Preprint (2024),
URL: https://scipost.org/preprints/scipost_202407_00017v1/.
- [102] D. Eberz, A. Kell, M. Breyer and M. Köhl,
Cooling a strongly-interacting quantum gas by interaction modulation, 2024, arXiv: 2410.10642 [cond-mat.quant-gas],
URL: <http://arxiv.org/abs/2410.10642>.
- [103] M. Köhl, *Atomic, Molecular, and Optical Physics*, 2017.
- [104] M. Randeria and E. Taylor, *Crossover from Bardeen-Cooper-Schrieffer to Bose-Einstein Condensation and the Unitary Fermi Gas*, Annual Review of Condensed Matter Physics **5** (2014) 209, URL: <https://www.annualreviews.org/doi/10.1146/annurev-conmatphys-031113-133829>.
- [105] A. Jabs, *Connecting Spin and Statistics in Quantum Mechanics*, Foundations of Physics **40** (2010) 776,
URL: <https://doi.org/10.1007/s10701-009-9351-4>.
- [106] C. Chin, R. Grimm, P. Julienne and E. Tiesinga,
Feshbach resonances in ultracold gases, Reviews of Modern Physics **82** (2010) 1225,
URL: <https://link.aps.org/doi/10.1103/RevModPhys.82.1225>.

- [107] Z.-C. Yan, J. F. Babb, A. Dalgarno and G. W. F. Drake, *Variational calculations of dispersion coefficients for interactions among H, He, and Li atoms*, Physical Review A **54** (1996) 2824,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.54.2824>.
- [108] S. Giorgini, L. P. Pitaevskii and S. Stringari, *Theory of ultracold atomic Fermi gases*, Reviews of Modern Physics **80** (2008) 1215,
URL: <https://link.aps.org/doi/10.1103/RevModPhys.80.1215>.
- [109] U. Fano, *Effects of Configuration Interaction on Intensities and Phase Shifts*, Physical Review **124** (1961) 1866,
URL: <https://link.aps.org/doi/10.1103/PhysRev.124.1866>.
- [110] T. Köhler, K. Góral and P. S. Julienne, *Production of cold molecules via magnetically tunable Feshbach resonances*, Reviews of Modern Physics **78** (2006) 1311,
URL: <https://link.aps.org/doi/10.1103/RevModPhys.78.1311>.
- [111] G. Zürn et al., *Precise Characterization of ^6Li Feshbach Resonances Using Trap-Sideband-Resolved RF Spectroscopy of Weakly Bound Molecules*, Physical Review Letters **110** (2013) 135301,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.110.135301>.
- [112] K. E. Strecker, G. B. Partridge and R. G. Hulet, *Conversion of an Atomic Fermi Gas to a Long-Lived Molecular Bose Gas*, Physical Review Letters **91** (2003) 080406,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.91.080406>.
- [113] J. Stenger et al., *Strongly Enhanced Inelastic Collisions in a Bose-Einstein Condensate near Feshbach Resonances*, Physical Review Letters **82** (1999) 2422,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.82.2422>.
- [114] B. Gao, *Solutions of the Schrödinger equation for an attractive $1/r^6$ potential*, Physical Review A **58** (1998) 1728,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.58.1728>.
- [115] B. Gao, *Binding energy and scattering length for diatomic systems*, Journal of Physics B: Atomic, Molecular and Optical Physics **37** (2004) 4273,
URL: <https://dx.doi.org/10.1088/0953-4075/37/21/004>.
- [116] G. F. Gribakin and V. V. Flambaum, *Calculation of the scattering length in atomic collisions using the semiclassical approximation*, Physical Review A **48** (1993) 546,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.48.546>.
- [117] V. N. Popov, *Theory of a Bose Gas Produced by Bound States of Fermi Particles*, JETP **23** (1966) 1034,
URL: <http://www.jetp.ras.ru/cgi-bin/e/index/e/23/6/p1034?a=list>.
- [118] L. V. Keldysh and A. N. Kozlov, *Collective properties of excitons in semiconductors*, JETP **27** (1968) 521,
URL: <http://www.jetp.ras.ru/cgi-bin/e/index/e/27/3/p521?a=list>.

-
- [119] D. M. Eagles, *Possible Pairing without Superconductivity at Low Carrier Concentrations in Bulk and Thin-Film Superconducting Semiconductors*, *Physical Review* **186** (1969) 456, URL: <https://link.aps.org/doi/10.1103/PhysRev.186.456>.
- [120] G. V. Skorniakov and K. A. Ter-Martirosian, *Three Body Problem for Short Range Forces. I. Scattering of Low Energy Neutrons by Deuterons*, *JETP* **4** (1957) 648, URL: <http://www.jetp.ras.ru/cgi-bin/e/index/e/4/5/p648?a=list>.
- [121] D. S. Petrov, C. Salomon and G. V. Shlyapnikov, *Weakly Bound Dimers of Fermionic Atoms*, *Physical Review Letters* **93** (2004) 090404, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.93.090404>.
- [122] G. Baym, J.-P. Blaizot, M. Holzmann, F. Laloë and D. Vautherin, *The Transition Temperature of the Dilute Interacting Bose Gas*, *Physical Review Letters* **83** (1999) 1703, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.83.1703>.
- [123] G. Baym, J.-P. Blaizot and J. Zinn-Justin, *The transition temperature of the dilute interacting Bose gas for N internal states*, *Europhysics Letters* **49** (2000) 150, URL: <https://epljournal.edpsciences.org/articles/epl/abs/2000/02/5927/5927.html>.
- [124] G. Baym, J.-P. Blaizot, M. Holzmann, F. Laloë and D. Vautherin, *Bose-Einstein transition in a dilute interacting gas*, *The European Physical Journal B* **24** (2001) 107, URL: <https://epjb.edpsciences.org/articles/epjb/abs/2001/21/b01370/b01370.html>.
- [125] H. Hu, X.-J. Liu and P. D. Drummond, *Equation of state of a superfluid Fermi gas in the BCS-BEC crossover*, *Europhysics Letters* **74** (2006) 574, URL: <https://epljournal.edpsciences.org/articles/epl/abs/2006/10/ep19344/ep19344.html>.
- [126] C. J. Pethick and H. Smith, *Bose-Einstein Condensation in Dilute Gases*, 2nd ed., Cambridge University Press, 2008, URL: <https://www.cambridge.org/core/books/boseeinstein-condensation-in-dilute-gases/CC439EAD70D78E47E9AF536DA7B203EC>.
- [127] M. Marini, F. Pistolesi and G. Strinati, *Evolution from BCS superconductivity to Bose condensation: analytic results for the crossover in three dimensions*, *The European Physical Journal B - Condensed Matter and Complex Systems* **1** (1998) 151, URL: <https://doi.org/10.1007/s100510050165>.
- [128] E. Burovski, E. Kozik, N. Prokof'ev, B. Svistunov and M. Troyer, *Critical Temperature Curve in BEC-BCS Crossover*, *Physical Review Letters* **101** (2008) 090402, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.101.090402>.
- [129] S. Hoinka et al., *Goldstone mode and pair-breaking excitations in atomic Fermi superfluids*, *Nature Physics* **13** (2017) 943, URL: <https://www.nature.com/articles/nphys4187>.

- [130] S. Tsuchiya, D. Yamamoto, R. Yoshii and M. Nitta, *Hidden charge-conjugation, parity, and time-reversal symmetries and massive Goldstone (Higgs) modes in superconductors*, Physical Review B **98** (2018) 094503, URL: <https://link.aps.org/doi/10.1103/PhysRevB.98.094503>.
- [131] V. Gurarie, *Nonequilibrium Dynamics of Weakly and Strongly Paired Superconductors*, Physical Review Letters **103** (2009) 075301, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.103.075301>.
- [132] R. Shimano and N. Tsuji, *Higgs Mode in Superconductors*, Annual Review of Condensed Matter Physics **11** (2020) 103, URL: <https://doi.org/10.1146/annurev-conmatphys-031119-050813>.
- [133] M. Link, *Exploring Non-Equilibrium in Ultracold Fermi Gases and Machine Learning in Physics*, PhD thesis: University of Bonn, 2020.
- [134] A. Kell, *Measuring Excitations of a Strongly Interacting Superfluid Fermi Gas*, PhD thesis: University of Bonn, 2022.
- [135] M. Breyer, *Higgs Mode and Critical Temperature in the BCS-BEC Crossover*, PhD thesis: University of Bonn, 2023.
- [136] M. E. Gehm, *Properties of ^6Li* , (2003), URL: <https://jet.physics.ncsu.edu/techdocs/>.
- [137] D. A. Steck, *Sodium D Line Data*, (2024), URL: <http://steck.us/alkalidata>.
- [138] C. Linse, *Construction of a ^6Li spectroscopy cell and Doppler-free spectroscopy*, BA thesis: University of Bonn, 2014.
- [139] E. D. Black, *An introduction to Pound–Drever–Hall laser frequency stabilization*, American Journal of Physics **69** (2001) 79, URL: <https://doi.org/10.1119/1.1286663>.
- [140] M. Zawierucha, *Design and Construction of a Sodium Spectroscopy Cell and a Lock-in Amplifier for Frequency Locking a Laser*, BA thesis: University of Bonn, 2015.
- [141] C. J. Foot, *Atomic physics*, Oxford master series in physics 7. Atomic, Optical, and laser physics, Oxford University Press, 2005.
- [142] C. A. Stan and W. Ketterle, *Multiple species atom source for laser-cooling experiments*, Review of Scientific Instruments **76** (2005) 063113, URL: <https://doi.org/10.1063/1.1935433>.
- [143] D. Eberz, *Design and Construction of a Multi-Species Effusive Atomic Oven for Lithium and Sodium*, MA thesis: University of Bonn, 2020.
- [144] S. C. Bell et al., *A slow atom source using a collimated effusive oven and a single-layer variable pitch coil Zeeman slower*, Review of Scientific Instruments **81** (2010) 013105, URL: <https://aip.scitation.org/doi/10.1063/1.3276712>.

-
- [145] Z. Hadzibabic et al., *Fiftyfold Improvement in the Number of Quantum Degenerate Fermionic Atoms*, Physical Review Letters **91** (2003) 160401, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.91.160401>.
- [146] W. Ketterle, K. B. Davis, M. A. Joffe, A. Martin and D. E. Pritchard, *High densities of cold atoms in a dark spontaneous-force optical trap*, Physical Review Letters **70** (1993) 2253, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.70.2253>.
- [147] J. Dalibard and C. Cohen-Tannoudji, *Laser cooling below the Doppler limit by polarization gradients: simple theoretical models*, JOSA B **6** (1989) 2023, URL: <https://opg.optica.org/josab/abstract.cfm?uri=josab-6-11-2023>.
- [148] P. D. Lett et al., *Observation of Atoms Laser Cooled below the Doppler Limit*, Physical Review Letters **61** (1988) 169, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.61.169>.
- [149] Z. Hadzibabic, *Studies of a Quantum Degenerate Fermionic Lithium Gas*, PhD thesis: MIT, 2003.
- [150] E. Majorana, *Atomi orientati in campo magnetico variabile*, Il Nuovo Cimento (1924-1942) **9** (1932) 43, URL: <https://doi.org/10.1007/BF02960953>.
- [151] C. V. Sukumar and D. M. Brink, *Spin-flip transitions in a magnetic trap*, Physical Review A **56** (1997) 2451, URL: <https://link.aps.org/doi/10.1103/PhysRevA.56.2451>.
- [152] D. M. Brink and C. V. Sukumar, *Majorana spin-flip transitions in a magnetic trap*, Physical Review A **74** (2006) 035401, URL: <https://link.aps.org/doi/10.1103/PhysRevA.74.035401>.
- [153] M.-S. Heo, J.-y. Choi and Y.-i. Shin, *Fast production of large ^{23}Na Bose-Einstein condensates in an optically plugged magnetic quadrupole trap*, Physical Review A **83** (2011) 013622, URL: <https://link.aps.org/doi/10.1103/PhysRevA.83.013622>.
- [154] J. Dalibard and C. Cohen-Tannoudji, *Dressed-atom approach to atomic motion in laser light: the dipole force revisited*, JOSA B **2** (1985) 1707, URL: <https://opg.optica.org/josab/abstract.cfm?uri=josab-2-11-1707>.
- [155] R. Grimm and M. Weidemüller, *Optical Dipole Traps for Neutral Atoms*, Academic Press, Volume 42, 2000, URL: <https://www.sciencedirect.com/science/article/abs/pii/S1049250X0860186X>.
- [156] L. Chomaz et al., *Emergence of coherence via transverse condensation in a uniform quasi-two-dimensional Bose gas*, Nature Communications **6** (2015) 6162, URL: <https://www.nature.com/articles/ncomms7162>.

- [157] N. Navon, R. P. Smith and Z. Hadzibabic, *Quantum Gases in Optical Boxes*, Nature Physics **17** (2021) 1334,
URL: <https://www.nature.com/articles/s41567-021-01403-z>.
- [158] Y. Wang et al.,
Preparation of hundreds of microscopic atomic ensembles in optical tweezer arrays, npj Quantum Information **6** (2020) 1,
URL: <https://www.nature.com/articles/s41534-020-0285-1>.
- [159] M. Gall, N. Wurz, J. Samland, C. F. Chan and M. Köhl,
Competing magnetic orders in a bilayer Hubbard model with ultracold atoms, Nature **589** (2021) 40,
URL: <https://www.nature.com/articles/s41586-020-03058-x>.
- [160] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics*,
Boston University: Wiley, 2007.
- [161] J. Andrijauskas, *Optical dipole trap laser setup in a new quantum gas experiment*,
MA thesis: University of Bonn, 2015.
- [162] C. Zener and R. H. Fowler, *Non-adiabatic crossing of energy levels*,
Proceedings of the Royal Society of London. Series A, Containing Papers of a
Mathematical and Physical Character **137** (1997) 696,
URL: <https://royalsocietypublishing.org/doi/10.1098/rspa.1932.0165>.
- [163] A. Kell et al., *A compact and fast magnetic coil for the manipulation of quantum
gases with Feshbach resonances*, Review of Scientific Instruments **92** (2021) 093202,
URL: <https://doi.org/10.1063/5.0049518>.
- [164] L. Montgomery Smith, D. R. Keefer and S. I. Sudharsanan,
Abel inversion using transform techniques,
Journal of Quantitative Spectroscopy and Radiative Transfer **39** (1988) 367, URL:
<https://www.sciencedirect.com/science/article/pii/002240738890101X>.
- [165] M. J. H. Ku, A. T. Sommer, L. W. Cheuk and M. W. Zwierlein,
*Revealing the Superfluid Lambda Transition in the Universal Thermodynamics of a
Unitary Fermi Gas*, Science **335** (2012) 563,
URL: <https://www.science.org/doi/10.1126/science.1214987>.
- [166] M. W. Zwierlein, C. H. Schunck, C. A. Stan, S. M. F. Raupach and W. Ketterle,
Formation Dynamics of a Fermion Pair Condensate,
Physical Review Letters **94** (2005) 180401,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.94.180401>.
- [167] P. A. Murthy et al.,
Matter-wave Fourier optics with a strongly interacting two-dimensional Fermi gas,
Physical Review A **90** (2014) 043611,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.90.043611>.
- [168] G. Reinaudi, T. Lahaye, Z. Wang and D. Guéry-Odelin,
Strong saturation absorption imaging of dense clouds of ultracold atoms,
Optics Letters **32** (2007) 3143,
URL: <https://www.osapublishing.org/abstract.cfm?URI=ol-32-21-3143>.

-
- [169] L. Chomaz, *Coherence and superfluidity of Bose gases in reduced dimensions : from harmonic traps to uniform fluids*, PhD thesis: Collège de France, 2014.
- [170] F. Schwabl, *Statistische Mechanik*, 3., aktualisierte Auflage, Springer, 2006.
- [171] Y. Hou and J. E. Drut,
Fourth- and Fifth-Order Virial Coefficients from Weak Coupling to Unitarity,
Physical Review Letters **125** (2020) 050403,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.050403>.
- [172] X.-J. Liu, H. Hu and P. D. Drummond,
Virial Expansion for a Strongly Correlated Fermi Gas,
Physical Review Letters **102** (2009) 160401,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.102.160401>.
- [173] D. Lee and T. Schäfer, *Cold dilute neutron matter on the lattice. I. Lattice virial coefficients and large scattering lengths*, Physical Review C **73** (2006) 015201,
URL: <https://link.aps.org/doi/10.1103/PhysRevC.73.015201>.
- [174] E. Beth and G. E. Uhlenbeck,
The quantum theory of the non-ideal gas. II. Behaviour at low temperatures,
Physica **4** (1937) 915, URL:
<https://www.sciencedirect.com/science/article/pii/S0031891437801895>.
- [175] X. Leyronas, *Virial expansion with Feynman diagrams*,
Physical Review A **84** (2011) 053633,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.84.053633>.
- [176] M. Holten et al.,
Observation of Cooper pairs in a mesoscopic two-dimensional Fermi gas,
Nature **606** (2022) 287,
URL: <https://www.nature.com/articles/s41586-022-04678-1>.
- [177] G. Carleo et al., *Machine learning and the physical sciences*,
Reviews of Modern Physics **91** (2019) 045002,
URL: <https://link.aps.org/doi/10.1103/RevModPhys.91.045002>.
- [178] A. Krizhevsky, I. Sutskever and G. E. Hinton,
ImageNet classification with deep convolutional neural networks,
Commun. ACM **60** (2017) 84, URL: <https://dl.acm.org/doi/10.1145/3065386>.
- [179] G. E. Hinton and R. R. Salakhutdinov,
Reducing the Dimensionality of Data with Neural Networks, Science **313** (2006) 504,
URL: <https://www.science.org/doi/10.1126/science.1127647>.
- [180] M. A. Nielsen, *Neural Networks and Deep Learning*, Publisher: Determination Press,
Determination Press, 2015, URL: <http://neuralnetworksanddeeplearning.com>.
- [181] I. Goodfellow, B. Yoshua and A. Courville, *Deep Learning*, MIT Press, 2016,
URL: <https://www.deeplearningbook.org/>.
- [182] G. Sanderson,
Neural Networks, The basics of neural networks, and the math behind how they learn,
(2024), URL: <https://www.3blue1brown.com/topics/neural-networks>.

- [183] F. Emmert-Streib, Z. Yang, H. Feng, S. Tripathi and M. Dehmer, *An Introductory Review of Deep Learning for Prediction Models With Big Data*, *Frontiers in Artificial Intelligence* **3** (2020), URL: <https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2020.00004/full>.
- [184] K. Hornik, M. Stinchcombe and H. White, *Multilayer feedforward networks are universal approximators*, *Neural Networks* **2** (1989) 359, URL: <https://www.sciencedirect.com/science/article/pii/0893608089900208>.
- [185] G. Cybenko, *Approximation by superpositions of a sigmoidal function*, *Mathematics of Control, Signals and Systems* **2** (1989) 303, URL: <https://doi.org/10.1007/BF02551274>.
- [186] F. Rosenblatt, *The perceptron - A perceiving and recognizing automaton*, tech. rep. 85-460-1, Cornell Aeronautical Laboratory, 1957.
- [187] X. Glorot, A. Bordes and Y. Bengio, *Deep Sparse Rectifier Neural Networks*, 2011, URL: <https://proceedings.mlr.press/v15/glorot11a.html>.
- [188] D. H. Hubel and T. N. Wiesel, *Receptive fields and functional architecture of monkey striate cortex*, *The Journal of Physiology* **195** (1968) 215, URL: <https://onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.1968.sp008455>.
- [189] Y. LeCun, Y. Bengio and G. Hinton, *Deep learning*, *Nature* **521** (2015) 436, URL: <https://www.nature.com/articles/nature14539>.
- [190] D. E. Rumelhart, G. E. Hinton and R. J. Williams, *Learning representations by back-propagating errors*, *Nature* **323** (1986) 533, URL: <https://www.nature.com/articles/323533a0>.
- [191] D. Masters and C. Luschi, *Revisiting Small Batch Training for Deep Neural Networks*, 2018, arXiv: 1804.07612 [cs.LG], URL: <http://arxiv.org/abs/1804.07612>.
- [192] D. P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, 2017, arXiv: 1412.6980 [cs.LG], URL: <http://arxiv.org/abs/1412.6980>.
- [193] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio and P.-A. Manzagol, *Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion*, *Journal of Machine Learning Research* **11** (2010) 3371, URL: <http://jmlr.org/papers/v11/vincent10a.html>.
- [194] O. Simeone, *A Very Brief Introduction to Machine Learning With Applications to Communication Systems*, 2018, arXiv: 1808.02342 [cs.IT], URL: <http://arxiv.org/abs/1808.02342>.
- [195] R. Rombach, A. Blattmann, D. Lorenz, P. Esser and B. Ommer, *High-Resolution Image Synthesis with Latent Diffusion Models*, 2022, arXiv: 2112.10752 [cs.CV], URL: <http://arxiv.org/abs/2112.10752>.

-
- [196] I. T. Jolliffe, *Principal Component Analysis*, Springer Series in Statistics, Springer, 1986, URL: <http://link.springer.com/10.1007/978-1-4757-1904-8>.
- [197] L. Alzubaidi et al., *Review of deep learning: concepts, CNN architectures, challenges, applications, future directions*, Journal of Big Data **8** (2021) 53, URL: <https://doi.org/10.1186/s40537-021-00444-8>.
- [198] K. Weiss, T. M. Khoshgoftaar and D. Wang, *A survey of transfer learning*, Journal of Big Data **3** (2016) 9, URL: <https://doi.org/10.1186/s40537-016-0043-6>.
- [199] C. Tan et al., *A Survey on Deep Transfer Learning*, 2018, arXiv: 1808.01974v1 [cs.LG], URL: <https://arxiv.org/abs/1808.01974v1>.
- [200] A. Shrikumar, P. Greenside and A. Kundaje, *Learning Important Features Through Propagating Activation Differences*, 2019, arXiv: 1704.02685 [cs.CV], URL: <http://arxiv.org/abs/1704.02685>.
- [201] G. Ness, A. Vainbaum, C. Shkedrov, Y. Florshaim and Y. Sagi, *Single-Exposure Absorption Imaging of Ultracold Atoms Using Deep Learning*, Physical Review Applied **14** (2020) 014011, URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.14.014011>.
- [202] G. Carleo and M. Troyer, *Solving the quantum many-body problem with artificial neural networks*, Science **355** (2017) 602, URL: <https://www.science.org/doi/10.1126/science.aag2302>.
- [203] G. Torlai et al., *Neural-network quantum state tomography*, Nature Physics **14** (2018) 447, URL: <https://www.nature.com/articles/s41567-018-0048-5>.
- [204] D.-L. Deng, X. Li and S. Das Sarma, *Machine learning topological states*, Physical Review B **96** (2017) 195145, URL: <https://link.aps.org/doi/10.1103/PhysRevB.96.195145>.
- [205] P. Zhang, H. Shen and H. Zhai, *Machine Learning Topological Invariants with Neural Networks*, Physical Review Letters **120** (2018) 066401, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.066401>.
- [206] E. Zhao et al., *Observing a topological phase transition with deep neural networks from experimental images of ultracold atoms*, 2022, arXiv: 2209.10060 [cond-mat.quant-gas], URL: <http://arxiv.org/abs/2209.10060>.
- [207] N. Käming, *Machine-Learning Enhanced Studies of Quantum Many-Body Systems using Experimental and Numerical Data*, PhD thesis: University of Hamburg, 2024, URL: <https://ediss.sub.uni-hamburg.de/handle/ediss/10759>.
- [208] J. Kinast et al., *Heat Capacity of a Strongly Interacting Fermi Gas*, Science **307** (2005) 1296, URL: <https://www.science.org/doi/10.1126/science.1109220>.

- [209] E. D. Kuhnle et al.,
Temperature Dependence of the Universal Contact Parameter in a Unitary Fermi Gas,
Physical Review Letters **106** (2011) 170402,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.106.170402>.
- [210] A. LeNail, *NN-SVG: Publication-Ready Neural Network Architecture Schematics*,
Journal of Open Source Software **4** (2019) 747,
URL: <https://joss.theoj.org/papers/10.21105/joss.00747>.
- [211] Martín Abadi et al.,
TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, 2015,
URL: <https://www.tensorflow.org/>.
- [212] L. H. Gilpin et al.,
Explaining Explanations: An Overview of Interpretability of Machine Learning, 2018,
URL: <https://ieeexplore.ieee.org/document/8631448>.
- [213] P. W. Anderson, *Plasmons, Gauge Invariance, and Mass*,
Physical Review **130** (1963) 439,
URL: <https://link.aps.org/doi/10.1103/PhysRev.130.439>.
- [214] P. B. Littlewood and C. M. Varma, *Amplitude collective modes in superconductors and their coupling to charge-density waves*, Physical Review B **26** (1982) 4883,
URL: <https://link.aps.org/doi/10.1103/PhysRevB.26.4883>.
- [215] S. Musolino,
Nonequilibrium dynamics of strongly correlated quantum gases: From few to many,
PhD thesis: Eindhoven University of Technology, 2021.
- [216] J. Tokimoto, S. Tsuchiya and T. Nikuni,
Excitation of the Higgs Mode in a Superfluid Fermi Gas in the BCS-BEC Crossover,
Journal of the Physical Society of Japan **88** (2019) 023601,
URL: <https://journals.jps.jp/doi/10.7566/JPSJ.88.023601>.
- [217] E. A. Yuzbashyan, M. Dzero, V. Gurarie and M. S. Foster,
Quantum quench phase diagrams of an s-wave BCS-BEC condensate,
Physical Review A **91** (2015) 033628,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.91.033628>.
- [218] M. Greiner, C. A. Regal and D. S. Jin,
Probing the Excitation Spectrum of a Fermi Gas in the BCS-BEC Crossover Regime,
Physical Review Letters **94** (2005) 070403,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.94.070403>.
- [219] C. Chin et al., *Observation of the Pairing Gap in a Strongly Interacting Fermi Gas*,
Science **305** (2004) 1128,
URL: <https://www.science.org/doi/10.1126/science.1100818>.
- [220] J. Plata, *Magnetic-modulation spectroscopy of an atomic Fermi gas in the BCS-BEC crossover: Dissociation spectra in the Bose-Einstein condensate regime*,
Physical Review A **74** (2006) 013603,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.74.013603>.

-
- [221] H. Kurkjian, S. N. Klimin, J. Tempere and Y. Castin,
Pair-Breaking Collective Branch in BCS Superconductors and Superfluid Fermi Gases,
Physical Review Letters **122** (2019) 093403,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.122.093403>.
- [222] A. Richie-Halford, J. E. Drut and A. Bulgac,
Emergence of a Pseudogap in the BCS-BEC Crossover,
Physical Review Letters **125** (2020) 060403,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.060403>.
- [223] J. P. Gaebler et al.,
Observation of pseudogap behaviour in a strongly interacting Fermi gas,
Nature Physics **6** (2010) 569, URL: <https://www.nature.com/articles/nphys1709>.
- [224] R. Combescot, A. Recati, C. Lobo and F. Chevy,
Normal State of Highly Polarized Fermi Gases: Simple Many-Body Approaches,
Physical Review Letters **98** (2007) 180402,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.98.180402>.
- [225] L. P. Pitaevskii and S. Stringari, *Bose-Einstein condensation and superfluidity*,
First edition, International series of monographs on physics 164,
Oxford, United Kingdom: Oxford University Press, 2016.
- [226] M. Popp, J.-J. Garcia-Ripoll, K. G. Vollbrecht and J. I. Cirac,
Ground-state cooling of atoms in optical lattices, Physical Review A **74** (2006) 013622,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.74.013622>.
- [227] B. Capogrosso-Sansone, Ş. G. Söyler, N. Prokof'ev and B. Svistunov,
Monte Carlo study of the two-dimensional Bose-Hubbard model,
Physical Review A **77** (2008) 015602,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.77.015602>.
- [228] J.-S. Bernier et al.,
Cooling fermionic atoms in optical lattices by shaping the confinement,
Physical Review A **79** (2009) 061601,
URL: <https://link.aps.org/doi/10.1103/PhysRevA.79.061601>.
- [229] C. S. Chiu, G. Ji, A. Mazurenko, D. Greif and M. Greiner,
Quantum State Engineering of a Hubbard System with Ultracold Fermions,
Physical Review Letters **120** (2018) 243201,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.243201>.
- [230] B. Yang et al., *Cooling and entangling ultracold atoms in optical lattices*,
Science **369** (2020) 550,
URL: <https://www.science.org/doi/10.1126/science.aaz6801>.
- [231] V. L. Berezinskii, *Destruction of Long-range Order in One-dimensional and Two-dimensional Systems having a Continuous Symmetry Group I. Classical Systems*,
JETP **32** (1971) 493,
URL: <http://www.jetp.ras.ru/cgi-bin/e/index/e/32/3/p493?a=list>.

- [232] P. Christodoulou et al., *Observation of first and second sound in a BKT superfluid*, Nature **594** (2021) 191,
URL: <https://www.nature.com/articles/s41586-021-03537-9>.
- [233] A. Raçon and N. Dupuis,
Higgs amplitude mode in the vicinity of a $(2 + 1)$ -dimensional quantum critical point,
Physical Review B **89** (2014) 180501,
URL: <https://link.aps.org/doi/10.1103/PhysRevB.89.180501>.
- [234] D. Phan and A. V. Chubukov,
Following the Higgs mode across the BCS-BEC crossover in two dimensions,
Physical Review B **107** (2023) 134519,
URL: <https://link.aps.org/doi/10.1103/PhysRevB.107.134519>.
- [235] E. J. Mueller, *Review of pseudogaps in strongly interacting Fermi gases*,
Reports on Progress in Physics **80** (2017) 104401,
URL: <https://iopscience.iop.org/article/10.1088/1361-6633/aa7e53>.
- [236] X. Li et al., *Observation and quantification of pseudogap in unitary Fermi gases*, 2023,
arXiv: 2310.14024 [cond-mat.quant-gas],
URL: <http://arxiv.org/abs/2310.14024>.
- [237] A. Kell, *Optical Trap for Ultracold 2D Fermi Gas*,
MA thesis: University of Bonn, 2017.
- [238] A. Franzen, *ComponentLibrary: a free vector graphics library for optics*,
URL: <http://www.gwoptics.org/ComponentLibrary/>.

Tools

This serves as a brief description of tools that have been used to create parts of this thesis.

- **Analysis:** Most parts of the analysis have been performed in *python*, mainly using standard libraries, such as *numpy*, *scipy*, *pandas* and *lmfit*. Data plotting has mainly employed the *matplotlib* library and the neural networks in this thesis have been created with *tensorflow*. Some calculations have been performed using *Wolfram Mathematica*.
- **CAD designs:** CAD designs have been created and rendered with a recent version of *Autodesk Inventor*.
- **Data extraction:** Data from third-party publications have been extracted using the *WebPlotDigitizer*, created by Ankit Rohatgi.
- **Thesis:** This thesis has been written in \LaTeX and the entries for this bibliography have been collected with *Zotero*.

Acknowledgements

Working on my PhD in the quantum group has been a great honour. First and foremost, I would like to express my gratitude to my supervisor, Prof. Dr. Michael Köhl. It has been invaluable to have a supervisor of such great expertise throughout my bachelor's, masters's and PhD studies. Thank you for giving me the opportunity to work on this remarkable experiment and with such a great team, which allowed us to produce many significant results.

Speaking of the team, I owe equal thanks to the *Humphry-Lab*. I joined the *Humphries* in July 2018 for my bachelor's thesis. At that point, the lab was composed of Martin Link, Andreas Kell and Kuiyi Gao, who welcomed me warmly and provided excellent advice and mentoring. Over the course of my PhD, many members joined and left, with a special mention to Moritz Breyer and Valentin Jonas. The PhD is a team effort, and I am truly grateful to all of you. A special thanks to Andreas Kell and Moritz Breyer, who accompanied me throughout the PhD journey. It has been a pleasure working with all of you.

The quantum group consists of more than just the *Humphry-Lab*, and many more contributed to this thesis. First, I would like to mention the other major legacy experiment, the *Fermi-Lab*, led by Nick Klemmer, Janek Fleper and Valentin Jonas. Further down the hall, another big experiment has been constructed, the *Fairy-Lab*, comprising Jonas Schmitz, Felix Rönchen, Tobias Kree, Ralf Berner, Tobias Grundmann and Thorsten Langerfeld. The newest addition to our group is the lab on *2D-Materials*, led by Moritz Scharfstädt and Andrea Bergschneider. Last but not least, I must mention the *Ion-lab*, ran by Santhosh Surendra, and the experiment on *cavity optomechanics* led by Steffen Gohlke. I will also not forget Akos Hoffmann, helping us out so many times, whether it was related to electronics, plumping or chemistry related issues.

Working with each of you has been a delight, and for the most part, it felt like working with friends rather than colleagues. A special thanks to everyone who helped us to wipe the floor after any of the several big *Humphry* water floods – last great flood being on 11th of May, 2024. My apologies to those of you (*Fairys*) effected by the strong magnetic fields of the *Humphry-Lab* – hopefully, you can forgive us.

My time in the group will always stick in my memory, and I would like to mention a few highlights: *Rosenkohl* in mensa, the corridor cupboard full of *Pfand* (we did great!), checking out *Massephase* on the scale in the *Bierverflüssigungsanlage*, *Oktoberfest* at Andreas von Haarens' place, going to *Ballermann* with Santhosh Surendra because it was raining too much in Bonn, *Surviving a Blizzard* in the mountains of Utah with Nick Klemmer, and off course the great gym gang.

As this thesis marks the *end* of the *Humphry-Lab*, I promised to declare my last will

regarding the fate of the lab equipment. While it remains to be seen if this will be honoured, it is *my wish* that Santhosh Surendra inherits all Delta power supplies and the Mephisto laser for his experiment.

I would also like to extend my thanks to the people working behind the scenes, specifically Tina Naggert and Dietmar Haubrich, for their great support and for their invaluable work maintaining the building's infrastructure.

Concluding on the people directly involved in this thesis, I would like to express special gratitude to those who took the time to proofread this thesis: Andreas Kell, Moritz Breyer, Valentin Jonas, Santhosh Surendra, Thorsten Langerfeld, Andreas von Haaren and Lea Lob. On this note, I would also like to thank the examiners for their time, namely Prof. Dr. Sebastian Hofferberth, Prof. Dr. David Luitz and Prof. Dr. Moritz Sokolowski.

Of course, my thanks extends to everyone who contributed to this thesis, even if not mentioned by name.

Gegen Ende dieser Danksagung möchte ich mich auch bei meiner Familie für ihre uneingeschränkte und bedingungslose Unterstützung bedanken. Es war eine große Erleichterung euch immer hinter mir zu wissen.

Abschließend möchte ich auch dir danke sagen, Lea, ich bin froh dich immer an meiner Seite zu wissen, und ich freue mich auf unsere gemeinsame Zukunft.